arXiv:1612.07717v1 [math.NA] 22 Dec 2016

Multilevel Monte Carlo and Improved Timestepping Methods in Atmospheric Dispersion Modelling

Grigoris Katsiolides[1], Eike H. Müller[*,1], Robert Scheichl[1], Tony Shardlow[1], Michael B. Giles[2] and David J. Thomson[3]

[1]*Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, United Kingdom*
[2]*Mathematical Institute, University of Oxford, Andrew Wiles Building, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, United Kingdom*
[3]*Met Office, Fitzroy Road, Exeter EX1 3PB, United Kingdom*
[*]*Email:* e.mueller@bath.ac.uk

December 23, 2016

Abstract

A common way to simulate the transport and spread of pollutants in the atmosphere is via stochastic Lagrangian dispersion models. Mathematically, these models describe turbulent transport processes with stochastic differential equations (SDEs). The computational bottleneck is the Monte Carlo algorithm, which simulates the motion of a large number of model particles in a turbulent velocity field; for each particle, a trajectory is calculated with a numerical timestepping method. Choosing an efficient numerical method is particularly important in operational emergency-response applications, such as tracking radioactive clouds from nuclear accidents or predicting the impact of volcanic ash clouds on international aviation, where accurate and timely predictions are essential. In this paper, we investigate the application of the Multilevel Monte Carlo (MLMC) method to simulate the propagation of particles in a representative one-dimensional dispersion scenario in the atmospheric boundary layer. MLMC can be shown to result in asymptotically superior computational complexity and reduced computational cost when compared to the Standard Monte Carlo (StMC) method, which is currently used in atmospheric dispersion modelling. To reduce the absolute cost of the method also in the non-asymptotic regime, it is equally important to choose the best possible numerical timestepping method on each level. To investigate this, we also compare the standard symplectic Euler method, which is used in many operational models, to an improved timestepping algorithm based on a geometric Langevin approach.

keywords: Atmospheric dispersion modelling, Multilevel Monte Carlo, Stochastic differential equations, Numerical timestepping methods

1 Introduction

Modelling the transport and dispersion of atmospheric pollutants is important in applied meteorological research, air-quality modelling for regulatory bodies and emergency-response applications. Numerical models have been used successfully to simulate the spread of radiactive clouds after the Fukushima nuclear accident in 2011 [1], and for evaluating the impact of volcanic ash on aviation following the 2010 and 2011 eruptions of the Eyjafjallajökull and Grímsvötn volcanoes on Iceland [2, 3]. The NAME atmospheric dispersion model [4, 5], originally designed to predict the spread of radiactive plumes after the Chernobyl disaster [6], is a highly versatile numerical model developed and used by the UK Met Office in a wide range

of applications, see e.g. [2, 3, 7]. Providing fast-yet-accurate and reliable predictions is crucial, in particular in emergency-response situations, where model outputs have to be passed on to decision makers in a timely fashion. At the heart of NAME, a Stochastic Differential Equation (SDE) describes the evolution of the distributions of position and velocity of the transported particles in the plume. This SDE is solved with a Monte Carlo method, which simulates the propagation of a large number of model particles in a turbulent velocity field. An efficient Monte Carlo algorithm and fast timestepping methods for integrating the trajectories of the individual particles are crucial for fast and precise forecasts. To assess the performance of a model, it is desirable to minimise the total computational cost for a given level of accuracy (measured by the root-mean-square error) in the prediction of a particular quantity of interest, such as the mean particle position or the concentration field at a given time.

Recently, the Multilevel Monte Carlo (MLMC) method [8, 9] (see [10] for a comprehensive review) has been used successfully in a wide range of applications from mathematical finance [11] to uncertainty quantification in subsurface flow [12]. Let ϵ denote the upper bound (or tolerance) on the root-mean-square error of the numerical method. This error consists of a deterministic discretisation error (bias) and a statistical error due to the limited number of model particles. To reduce both error components, the Standard Monte Carlo (StMC) method usually requires an asymptotic cost $\propto N \cdot M$ since N model particles are propagated over M timesteps. For a first-order method, the number of timesteps (for a given simulation duration) is $M \propto \epsilon^{-1}$ and the central limit theorem implies $N \propto \epsilon^{-2}$, since the statistical error decreases $\propto 1/\sqrt{N}$. In total, this results in an asymptotic cost $\text{Cost}_{\text{StMC}} \leq C_{\text{StMC}} \cdot \epsilon^{-3}$. In contrast, the MLMC algorithm can achieve significantly better scaling $\text{Cost}_{\text{MLMC}} \leq C_{\text{MLMC}} \cdot \epsilon^{-2}$ by introducing a hierarchy of additional, coarser levels with $M/2, M/4, \ldots$ timesteps. On each pair of subsequent levels, paths with two different timestep sizes but the same random noise are used to calculate the same approximation to a particular idealised trajectory. By summing up the contributions on all levels with a telescoping sum, the expected value on the original finest level with M timesteps is recovered. This reduces the overall cost for two reasons. Firstly, the cost of one path calculation on a coarser level is proportional to the (smaller) number of timesteps on this level. In addition, only differences between subsequent levels are computed. Those differences have a much smaller variance and can be evaluated with a smaller number of samples. The overall reduction in computational cost is quantified in more detail in the complexity theorem from [9], which is reviewed later in this paper. The absolute runtime also depends on the constants C_{StMC} and C_{MLMC}, which are very sensitive to the numerical timestepping method. This has recently been demonstrated for a set of simpler model problems in [13]. The relative size of those constants is particularly important for relatively loose tolerances, which are typical in operational atmospheric dispersion applications.

Building on [13], the aim of this work is to investigate the performance enhancements that can be achieved with MLMC and improved timestepping algorithms in atmospheric modelling. This is a challenging problem for several reasons: Firstly, the SDE is nonlinear and the coefficients vary strongly or even diverge in parts of the domain. When simulating vertical dispersion in a neutrally stable atmospheric boundary layer, the "profiles" for the turbulent velocity fluctuations σ_U and velocity decorrelation time τ have singularities near the ground and at the top of the domain and need to be regularised appropriately. In addition, suitable boundary conditions have to be applied to ensure that the model particles do not leave the atmospheric boundary layer. Finally, commonly used quantities of interest (in particular the particle concentration) are expressed as the expected value of a function which is not Lipschitz continuous and needs to be approximated correctly. All those issues are addressed in this paper.

While reflective boundary conditions for a one-dimensional diffusion process of the form $dX = a\,dt + \sigma\,dW$ have previously been treated by adding a non-decreasing "reflection" function (local time) to the particle position (see e.g. [14]), this method can not be directly applied to the system of two coupled SDEs for velocity and position that we study here. Instead, we use an approach similar to the *impact law* described for piecewise-smooth dynamical systems in [15]. More specifically, to treat reflective boundary conditions for SDEs, we follow the work in [16], who give an algorithm for elastic reflection at the top and bottom of the domain for a one-dimensional (vertical) dispersion model. The physical justification and implementation of various boundary conditions for atmospheric dispersion models is further discussed in [17, Chap. 11] and the references given there. Care has to be taken when using reflective boundary conditions in a MLMC context since trajectories of particles on subsequent levels might diverge. To ensure the strong convergence

of paths on subsequent levels, we consider an equivalent problem without boundary conditions, which can be obtained by periodically extending the domain. This new approach only requires the inclusion of factors of ± 1 in the coupling between random numbers on the coarse and fine levels.

To predict particle concentrations, we replace the expected value of the indicator function by a suitable smoothing polynomial, matching the first r moments of the indicator function. As discussed in [18], any additional errors introduced by this process can be systematically controlled and we confirm numerically that they are of the same size as the bias and statistical error from the Monte Carlo method.

Several advanced timestepping methods for SDEs have been discussed in the literature (see e.g. [19, 20]). As a reference method, we use the simple symplectic Euler integrator, which is implemented in many operational atmospheric dispersion models. This integrator is very similar to the well-known Euler–Maruyama method [21, 19] and has the same restrictions on the maximal timestep size, which become more severe at the bottom of the boundary layer due to the small velocity decorrelation time. To avoid those stability issues, we also investigate the performance of an improved timestepping method based on geometric Langevin methods [22]. This method had previously been applied in the context of MLMC methods for simpler, linear model problems in [13]. Similar splitting methods for Langevin dynamics are also discussed in the context of molecular dynamics in [20].

Our numerical results confirm that, for tight tolerances, due to its asymptotic scaling with ϵ^{-2} the MLMC algorithm always outperforms a standard Monte Carlo method. This difference in performance is particularly pronounced if the particles spend less time close to the ground where the profiles τ and σ_U have large variations. For looser tolerances, the standard Monte Carlo algorithm can be faster than MLMC. However, it is very sensitive to the timestepping method and using the geometric Langevin method can result in speedups of more than one order of magnitude, in particular if particles are released close to the ground.

An alternative way to avoid stability issues due to small velocity decorrelation times at the bottom of the boundary layer is adaptive timestepping. Two adaptive MLMC algorithms are described and analysed in [23, 24]. The authors of [24] either use a sample averaged error to construct a mesh hierarchy, which is used for all realisations, or adapt the timestep size for each path independently, again via an error estimator. In this work, we focus on stability issues and therefore employ a simple heuristic based on the local velocity decorrelation time. The implementation of the adaptive MLMC algorithm that we use is described in [25]. We find that adaptive timestepping can further reduce the cost for the symplectic Euler integrator. However, for the geometric Langevin method, our numerical experiments seem to suggest that uniform timestepping methods lead to a shorter overall solution time.

To obtain realistic absolute performance numbers, all results in this paper were obtained with a freely available object-oriented C++ code [26], which can be downloaded under the GPL license.

Structure This paper is organised as follows: in Section 2, we introduce the stochastic model for simulating particle transport and dispersion in the stationary inhomogeneous turbulence in the boundary layer and describe the underlying SDE for particle transport and dispersion. Section 3 introduces the timestepping methods, in particular the geometric Langevin integrator, and contains a review of the MLMC algorithm with particular focus on its computational complexity. The proper treatment of reflective boundary conditions in the context of MLMC is discussed in Section 4 and concentration estimates with suitable smoothed indicator functions are analysed in Section 5. In Section 6, we present numerical results for three quantities of interest: the mean particle position, the concentration at a particular location and the probability density function of particle position. This includes comparisons between the StMC and the MLMC methods so as to identify the most efficient combination. The impact of the regularisation height on the results and adaptive timestepping methods are discussed in Section 7. We finally conclude and outline ideas for further work in Section 8. Some more technical details are relegated to the appendices, where we analyse our methods based on the theory of modified equations [27, 28] and present some mathematical derivations related to the approximation of indicator (box) functions with smoothing polynomials.

2 Lagrangian Models for Atmospheric Dispersion

Dispersion models predict the spread and transport of atmospheric pollutants in a given velocity field, i.e. the dispersed material is treated as a passive non-interacting tracer. Usually, this velocity field is obtained from the output of a numerical weather prediction (NWP) model. Since the NWP model has a finite grid resolution, the total velocity field $u(x,t)$ can be split into two components. The evolution of the large-scale flow, which is resolved by the grid, is described by the Navier–Stokes equations and simulated as a deterministic velocity field $v(x,t)$ in the dynamical core of the model. In contrast, unresolved turbulence is described by an additional component $U(x,t)$, which is not known in detail and whose properties and effects can only be described statistically. In general, the velocity can be decomposed as

$$u(x,t) = v(x,t) + U(x,t), \qquad \text{where} \qquad \langle u(x,t)\rangle = v(x,t) \tag{1}$$

and angular brackets $\langle\cdot\rangle$ denote the ensemble average. For an observable quantity $A(x,t)$, the ensemble average is defined by

$$\langle A(x,t)\rangle = \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \alpha^{(n)}(x,t), \tag{2}$$

where $\alpha^{(n)}(x,t)$ is the n^{th} realisation of A, i.e. particle transport in one particular realisation of the turbulent velocity component $U(x,t)$.

Lagrangian dispersion models describe the motion of atmospheric tracer particles in turbulent atmospheric flow fields by a stochastic differential equation (SDE): each particle follows the mean flow $v(x,t)$ and turbulence is simulated by adding a random velocity component U_t. At equilibrium, the Lagrangian velocity U_t of particles which pass through point x at time t agrees in distribution with the turbulent background velocity field $U(x,t)$. For simplicity, we assume that the background flow field vanishes and will set $v(x,t) = 0$ in the following.

In [29], several criteria for the selection of models for simulating particle trajectories in turbulent flows are discussed (see also [17] for an overview). For simplicity, we assume here that turbulence in each space direction can be treated independently and only consider the vertical component of the turbulent motion in this article; studying the more general three-dimensional case is the topic of further research. The following one-dimensional model for particle dispersion in a stationary inhomogeneous turbulent flow is derived in [29, 17]:

$$dU_t = -\frac{U_t}{\tau(X_t)}\,dt + \frac{1}{2}\left[1 + \left(\frac{U_t}{\sigma_U(X_t)}\right)^2\right] \frac{\partial\sigma_U^2(X_t)}{\partial X}\,dt + \sqrt{\frac{2\sigma_U^2(X_t)}{\tau(X_t)}}\,dW_t, \qquad dX_t = U_t\,dt. \tag{3}$$

Here U_t is the Lagrangian vertical velocity fluctuation, X_t is the particle's position and W_t is a Brownian motion. To complete the description in Eq. (3), initial conditions for U_t and X_t at time $t = 0$ need to be specified. For stationary turbulence, $U(x,t)$ and $U(x,t_0)$ agree in distribution for any two times t and t_0. The velocity variance is $\sigma_U^2(x) = \langle U^2(x,t_0)\rangle$ where t_0 is an arbitrary reference time. Using Kolmogorov similarity theory [30], the Lagrangian local velocity decorrelation timescale τ can be expressed in terms of the mean rate of dissipation of turbulent kinetic energy ϵ_{turb} as $\tau = 2\sigma_U^2/(C_0\epsilon_{\text{turb}})$ [17]. Although for homogeneous turbulence, τ is identical to the Lagrangian velocity autocorrelation time of U_t, this is not true in the more general case of inhomogeneous turbulence considered here.

The term proportional to $\partial\sigma_U^2(X_t)/\partial X$ in Eq. (3) arises from the "well-mixed condition" [29], which prevents the unphysical spatial agglomeration of particles in an initially homogeneous concentration field. The well-mixed condition requires the equilibrium (X_t, U_t)-distribution of particles in the model to be the same as the distribution of all air parcels. Mathematically this is equivalent to requiring that the equilibrium distribution or invariant measure $p(X_t, U_t)$ is uniform in space and, conditional on $X_t = x$, is the same as the distribution of $U(x,t)$. The latter is independent of t because of our stationarity assumption and is here assumed to be Gaussian i.e. $p(U_t|X_t = x) = \frac{1}{\sqrt{2\pi}\sigma_U(x)} \exp(-U_t^2/2\sigma_U^2(x))$.

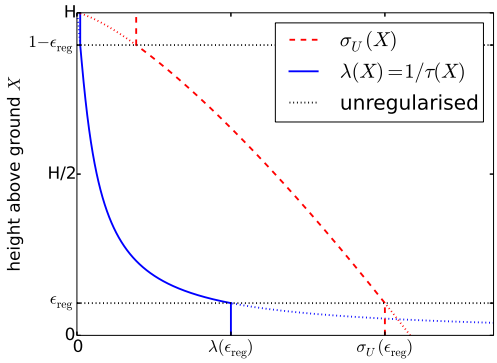

Figure 1: Vertical dependency of the profiles $\sigma_U(X)$ and $\lambda(X)$.

To simplify notation we introduce the following abbreviations

$$\lambda(X) \equiv \frac{1}{\tau(X)}, \qquad \sigma(X) \equiv \sqrt{\frac{2\sigma_U^2(X)}{\tau(X)}}, \qquad \frac{\partial V(X,U)}{\partial X} \equiv -\frac{1}{2}\left[1+\left(\frac{U}{\sigma_U(X)}\right)^2\right]\frac{\partial \sigma_U^2(X)}{\partial X}$$

so that Eq. (3) can be written more compactly as

$$dU_t = -\lambda(X_t)\,U_t\,dt - \frac{\partial V(X_t,U_t)}{\partial X_t}\,dt + \sigma(X_t)\,dW_t, \qquad dX_t = U_t\,dt \tag{4}$$

To complete this model, suitable vertical profiles for $\sigma_U(X)$ and $\tau(X)$ have to be chosen. For a neutrally stable surface layer and horizontally homogeneous turbulence, the authors of [31] derive the expression $\tau(X) \approx 0.5X/\sigma_U$ for the velocity decorrelation time. If the turbulent eddy size is only constrained by the ground, then a typical eddy size is $\propto X$. Hence, this relationship expresses the fact that on average it takes a time $\propto X/\sigma_U$ for a particle to decorrelate by going around one eddy. As we will see below, the fact that the decorrelation timescale $\tau(X)$ vanishes linearly for $X \to 0$ has important implications on the performance of the model. For the surface layer, the authors of [31] also suggest $\sigma_U = 1.3u_*$ where u_* is the friction velocity. There is considerable uncertainty in the exact form of the profiles further away from the surface. Rather than exploring a wide range of different parametrisations, we concentrate on one representative case, which allows us to quantify the impact of variations in σ_U on the algorithms discussed here. In this paper, we use the form of σ_U suggested in [32] for neutral and stable conditions. In summary, the following profiles are used to describe the propagation of particles in a boundary layer of depth H:

$$\sigma_U(X) = \kappa_\sigma u_* \left(1-\frac{X}{H}\right)^{\frac{3}{4}}, \qquad \tau(X) = \kappa_\tau \frac{X}{\sigma_U(X)}, \qquad X \in [0,H], \tag{5}$$

and the constants are set to $\kappa_\sigma = 1.3$, $\kappa_\tau = 0.5$, $u_* = 0.2\,\mathrm{m/s}$. Those profiles are shown in Fig. 1; as discussed in Section 2.1, τ (or λ) needs to be regularised near the bottom boundary. Since σ_U and τ are only defined within the boundary layer, suitable boundary conditions that restrict the particles to the interval $[0,H]$ need to be enforced.

Different boundary conditions for atmospheric dispersion models are discussed in [17]. Here we assume that when a particle hits the boundary, it is reflected elastically (see Fig. 2a). Mathematically this can be formulated as in [15]. If at some time $t^{(\mathrm{refl})}$ a particle hits the lower ($X = 0$) or the upper ($X = H$) boundary, we use

$$\lim_{h\to 0, h>0} U_{t^{(\mathrm{refl})}+h} = -\lim_{h\to 0, h>0} U_{t^{(\mathrm{refl})}-h}. \tag{6}$$

With initial conditions at $t = 0$, it should be noted that, under the assumption that $\tau(X)$ is bounded from below by a non-zero positive constant and that $\sigma_U(X)$ and $\tau(X)$ are sufficiently smooth, it is possible to show existence and uniqueness of a solution for the inhomogeneous model for all $t > 0$. The key idea is to

find a suitable Lyapunov function and then follow the theory developed in [33]; including the details of the proof is beyond the scope of this paper.

Given Eqs. (4) and an initial condition, it is now possible to calculate any functional of the path U_t, X_t. Any such functional $\phi(U_t, X_t)$ is referred to as a quantity of interest (QoI) in the following. Of particular interest in practical applications is the concentration field at a given time, which can be approximated by the probability of X_t lying in one of the boxes of a predetermined grid. Mathematically for a box B, this probability can be expressed as the expected value of the indicator function $\mathbb{1}_B(x)$ which is 1 for $x \in B$ and 0 otherwise. Indeed, such piecewise averaged concentration fields are one of the main outputs of widely used dispersion models such as NAME.

In this paper, we concentrate on calculating the box-averaged concentration $\mathbb{E}\left[\mathbb{1}_{[a,b]}(X_T)\right]$ for some $T > 0$, i.e. the fraction of particles that land in a specified interval $[a, b] \subset [0, H]$. As we will show in Section 6.3, it is possible to generalise this to calculate the piecewise constant concentration field at the final time by considering a vector of indicator functions. Since it is easier to predict, as a second quantity of interest, we will also consider the mean particle position $\mathbb{E}[X_T]$.

2.1 Regularisation

In Eq. (3), the profiles $\sigma_U(X)$ and $\tau(X)$ have singularities or vanish at the boundaries $X = 0$ and $X = H$. Consequently, $\partial V(X, U)/\partial X$ diverges for $X \to H$ and $\lambda(X) \to \infty$ for $X \to 0$. The latter restricts the timestep in numerical methods to an infinitely small value. It turns out that the issue is not merely the large value of λ, which could be treated by integrating a larger part of the equation exactly. The function $\lambda(X)$ also has an infinitely large gradient near the ground, which leads to more serious stability issues. In order to deal with these singularities, the functions $\tau(X)$ and $\sigma_U(X)$ have to be regularised. For this, we use the same approach as the Met Office NAME model by imposing a hard cutoff

$$\tau(X) = \begin{cases} \tau(\epsilon_{\text{reg}}) & \text{if} \quad X < \epsilon_{\text{reg}}, \\ \tau(H - \epsilon_{\text{reg}}) & \text{if} \quad X > H - \epsilon_{\text{reg}}, \end{cases} \qquad \sigma_U(X) = \begin{cases} \sigma_U(\epsilon_{\text{reg}}) & \text{if} \quad X < \epsilon_{\text{reg}}, \\ \sigma_U(H - \epsilon_{\text{reg}}) & \text{if} \quad X > H - \epsilon_{\text{reg}}, \end{cases} \tag{7}$$

where $\epsilon_{\text{reg}} \ll H$ is some small regularisation constant. This is shown for $\lambda(X) = 1/\tau(X)$ in Fig. 1. We will study the impact of ϵ_{reg} on the results in Section 7.1.

3 Algorithms

3.1 Timestepping methods

Timestepping methods for SDEs and especially the Langevin equation are described in [19, 13, 20]. To solve the SDE (4) numerically, we use two explicit timestepping methods that are both of weak order one for smooth problems (i.e. the bias error for fixed travel time grows linearly with the timestep size). The simulated time interval $[0, T]$ is split into M intervals of size $h = T/M$ and, for any $n \in 0, \ldots, M$, the quantity (U_n, X_n) denotes the numerical approximation of $(U(t_n), X(t_n))$ at time $t_n = nh$.

The first timestepping method is the Symplectic Euler method given by

$$\begin{aligned} U_{n+1} &= (1 - \lambda(X_n)\, h)\, U_n - \frac{\partial V(X_n, U_n)}{\partial X}\, h + \sigma(X_n)\sqrt{h}\, \xi_n, \\ X_{n+1} &= X_n + U_{n+1}\, h. \end{aligned} \tag{8}$$

For each timestep n, the random number $\xi_n \sim \text{Normal}(0, 1)$, independently of all previous timesteps. Due to its simplicity it is very popular in operational models, such as the Met Office's NAME code, and it will

serve as the benchmark in the following. It is referred to as "symplectic" because, when applied to a Hamiltonian system of ordinary differential equations, it inherits a symplectic structure from the Hamiltonian system.

The method in (8) is only stable if $|1 - \lambda(X_n)h| < 1$, which leads to the constraint $h < 2/\lambda(X_n)$ on the timestep size. For large values of $\lambda(X)$, which occur close to the lower boundary $X \to 0$, the timestep size has to be reduced and this increases the computational cost of the method. This problem can be mitigated in adaptive timestepping algorithms, which will be discussed in Section 7.

The geometric Langevin method [22] defined by

$$\begin{aligned} U^\star_{n+1} &= e^{-\lambda(X_n)h}\, U_n + \sigma(X_n)\sqrt{\frac{1 - e^{-2\lambda(X_n)h}}{2\lambda(X_n)}}\,\xi_n, \\ U_{n+1} &= U^\star_{n+1} - \frac{\partial V(X_n, U^\star_{n+1})}{\partial X}\,h, \\ X_{n+1} &= X_n + U_{n+1}\,h \end{aligned} \tag{9}$$

avoids this stability issue. As above, the random variables $\xi_n \sim \mathrm{Normal}(0,1)$ i.i.d. This method was introduced in the context of MLMC in [13] and referred to as "Symplectic Euler/Ornstein–Uhlenbeck Splitting Method" since, for constant σ_U and λ, the U equations are solved exactly as the Ornstein–Uhlenbeck process $dU_t = -\lambda\, U_t\, dt + \sigma\, dW_t$. The formulation in Eq. (9) is slightly more general in that the potential function V also depends on the velocity U.

3.2 Standard Monte Carlo simulation

In a standard Monte Carlo (StMC) simulation, N independent trajectories $\{(U_n^{(i)}, X_n^{(i)}), i = 1, \ldots, N\}$ are computed, each associated with a particular i.i.d. realisation of the Gaussian increments $\xi_n^{(i)}$, $n = 0, \ldots, M$, in (8) or (9). A quantity of interest (QoI) is evaluated for each trajectory. Possible QoIs are the final particle position or a variable which is one if the particle falls into a defined interval and zero otherwise. The expected value of the QoI can then be estimated by averaging the N realisations of the QoI. In the first example the expectation value is the mean particle position; in the second case it is the fraction of particles which fall into the chosen interval.

We will only consider functionals of the solution X_T at the final time T as QoIs to analyse the cost of the StMC method. In particular, let $\mathcal{P} = \phi(X_T)$. Given integer constants L and M_0, we approximate $\mathcal{P}$ by $\mathcal{P}_L = \phi(X_{M_L})$ where X_{M_L} is the numerical approximation of the particle position X_t at the final time $T = M_L h_L$, obtained with $M = M_L = M_0 2^L$ timesteps of size $h_L = \frac{T}{M_L} = \frac{T}{M_0} 2^{-L}$. The StMC estimator with $N = N_L$ samples for $\mathbb{E}[\mathcal{P}]$ is given by

$$\hat{\mathcal{P}}_L^{(\mathrm{StMC})} = \frac{1}{N_L}\sum_{i=1}^{N_L} \mathcal{P}_L^{(i)}, \qquad \text{where} \quad \mathcal{P}_L^{(i)} = \phi(X_{M_L}^{(i)}). \tag{10}$$

Due to the linearity of the expectation, $\mathbb{E}\left[\hat{\mathcal{P}}_L^{(\mathrm{StMC})}\right] = \mathbb{E}\left[\mathcal{P}_L^{(i)}\right] = \mathbb{E}\left[\mathcal{P}_L\right]$. It is also easy to show that the total mean-squared error of the StMC estimator in Eq. (10) is given by

$$\mathbb{E}\left[\left(\hat{\mathcal{P}}_L^{(\mathrm{StMC})} - \mathbb{E}\left[\mathcal{P}\right]\right)^2\right] = \left(\mathbb{E}\left[\mathcal{P}_L\right] - \mathbb{E}\left[\mathcal{P}\right]\right)^2 + \frac{1}{N_L}\mathrm{Var}\left[\mathcal{P}_L\right]. \tag{11}$$

The first term on the right hand side of Eq. (11) is the square of the numerical bias and, for the first-order methods described in Eqs. (8) and (9) above, this bias grows linearly with the timestep size: $\mathbb{E}\left[\mathcal{P}_L - \mathcal{P}\right] \propto h_L \propto M_L^{-1}$. The second term is the (mean square) statistical error, which is proportional to N_L^{-1}. To limit the total mean-squared error to ϵ^2 (where ϵ is a fixed tolerance), it is sufficient to choose M_L and N_L such that each of the two terms on the right hand side of Eq. (11) is smaller than $\epsilon^2/2$. As already discussed

in the introduction, this choice leads to $N_L \propto \epsilon^{-2}$ and $M_L \propto \epsilon^{-1}$, resulting in a total cost of the StMC method

$$\text{Cost}^{(\text{StMC})} \leq C_{\text{StMC}} \cdot \epsilon^{-3} \propto N_L \cdot M_L. \tag{12}$$

The size of the constant C_{StMC} depends on the timestepping method. As we will argue in the following, the MLMC approach can lead to improved asymptotic scaling with asymptoptic cost proportional to ϵ^{-2}, which is the optimal computational complexity for a Monte Carlo method.

3.3 Multilevel Monte Carlo simulation

We generalise the notation from the previous section by considering an arbitrary level $\ell \leq L$ with $M_\ell = M_0 2^\ell$ timesteps of size $h_\ell = \frac{T}{M_\ell} = \frac{T}{M_0} 2^{-\ell}$ and N_ℓ Monte Carlo samples. On each level, we define the random variable $\mathcal{P}_\ell = \phi(X_{M_\ell})$ where X_{M_ℓ} is the approximation of the particle position X_T at time $T = M_\ell h_\ell$, now calculated with a timestep size h_ℓ. $\mathcal{P}_\ell^{(i)} = \phi(X_{M_\ell}^{(i)})$ is the value of this random variable for the i-the sample.

The key idea for the MLMC approximation [8, 9, 10] is that $\mathbb{E}\left[\mathcal{P}_L\right]$ can be written as the following telescoping sum:

$$\mathbb{E}\left[\mathcal{P}_L\right] = \sum_{\ell=1}^{L} \mathbb{E}\left[\mathcal{P}_\ell - \mathcal{P}_{\ell-1}\right] + \mathbb{E}\left[\mathcal{P}_0\right]. \tag{13}$$

The MLMC estimator is then defined by

$$\hat{\mathcal{P}}_L^{(\text{MLMC})} = \sum_{\ell=0}^{L} \hat{Y}_{\ell,N_\ell}, \tag{14}$$

where

$$\hat{Y}_{0,N_0} = \hat{\mathcal{P}}_0^{(\text{StMC})} \quad \text{and} \quad \hat{Y}_{\ell,N_\ell} = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} Y_\ell^{(i)} \quad \text{with} \quad Y_\ell^{(i)} = \mathcal{P}_\ell^{(i)} - \mathcal{P}_{\ell-1}^{(i)}, \quad \text{for} \quad \ell \geq 1. \tag{15}$$

Since the expected value is linear, the MLMC estimator does not introduce any additional bias and

$$\mathbb{E}\left[\hat{\mathcal{P}}_L^{(\text{MLMC})}\right] = \sum_{\ell=1}^{L} \mathbb{E}\left[\mathcal{P}_\ell - \mathcal{P}_{\ell-1}\right] + \mathbb{E}\left[\mathcal{P}_0\right] = \mathbb{E}\left[\mathcal{P}_L\right] = \mathbb{E}\left[\hat{\mathcal{P}}_L^{(\text{StMC})}\right].$$

The estimators $\hat{Y}_{\ell,N_\ell}$ on two subsequent levels ℓ and $\ell+1$ are independent: even though they are evaluated with the same timestep size, the samples $\mathcal{P}_\ell^{(i)}$ required to calculate $\hat{Y}_{\ell,N_\ell}$ and $\hat{Y}_{\ell+1,N_{\ell+1}}$ in Eq. (15) are generated with different random numbers. Due to this independence of the estimators $\hat{Y}_{\ell,N_\ell}$ the mean-square error of $\hat{\mathcal{P}}_L^{(\text{MLMC})}$ is given by

$$\mathbb{E}\left[\left(\hat{\mathcal{P}}_L^{(\text{MLMC})} - \mathbb{E}\left[\mathcal{P}\right]\right)^2\right] = \left(\mathbb{E}\left[\mathcal{P}_L\right] - \mathbb{E}\left[\mathcal{P}\right]\right)^2 + \sum_{\ell=0}^{L} \frac{1}{N_\ell} \text{Var}\left[Y_\ell\right], \tag{16}$$

which should be compared to the corresponding expression for StMC in Eq. (11). To calculate $\hat{Y}_{\ell,N_\ell}$ on each level $\ell > 0$ and to evaluate the MLMC estimator in Eq. (14), N_ℓ samples of random numbers $\xi_n^{(\ell,i)}$, $i = 1, \ldots, N_\ell$, for each timestep $n = 1, \ldots, M_\ell$ are generated. For the ith sample, both a path with M_ℓ timesteps and a path with $M_\ell/2$ timesteps are calculated for the *same* driving Brownian motion W_t. The difference $Y_\ell^{(i)}$ between the QoI on the two paths is averaged according to Eq. (15). On the coarsest level $\ell = 0$, the standard MC estimator with N_0 samples and M_0 timesteps is used.

The two main advantages of the MLMC algorithm are:

- Since $\ell < L$ and hence $M_\ell < M_L$, paths on the coarse levels require fewer timesteps and are cheaper to evaluate. In the case we consider here, it can be shown that the computational cost is concentrated on the coarsest level with $M_0 = \mathcal{O}(1)$ timesteps, independent of L and of the fine level timestep size.
- Instead of calculating the quantity of interest itself, only differences $Y_\ell = \mathcal{P}_\ell - \mathcal{P}_{\ell-1}$ of the QoI between subsequent levels are evaluated. Those differences have a smaller variance than $\mathcal{P}_\ell$. As can be seen from Eq. (16), this has a direct impact on the total error.

As discussed in detail in [9], for a given root-mean-square error ϵ, it is now possible to choose the N_ℓ such that the total cost $\propto \sum_{\ell=0}^{L} N_\ell M_\ell$ is minimised. This is quantified in the following simplified version of the complexity theorem proved in [9].

Complexity Theorem. *With the MLMC setting as above and if there exist estimators $\hat{Y}_{\ell,N_\ell}$ and positive constants c_1, c_2, c_3 such that*

(i) $\left|\mathbb{E}[\hat{\mathcal{P}}_\ell - \mathcal{P}]\right| \leq c_1 h_\ell,$

(ii) $\mathbb{E}[\hat{Y}_{\ell,N_\ell}] = \begin{cases} \mathbb{E}[\mathcal{P}_0] & \text{if} \quad \ell = 0, \\ \mathbb{E}[\mathcal{P}_\ell - \mathcal{P}_{\ell-1}] & \text{if} \quad \ell > 0, \end{cases}$

(iii) $\mathrm{Var}[\hat{Y}_{\ell,N_\ell}] \leq c_2 N_\ell^{-1} h_\ell^2,$

(iv) C_ℓ*, the computational complexity of* $\hat{Y}_{\ell,N_\ell}$*, is bounded by* $c_3 N_\ell h_\ell^{-1}$*,*

then there exists a positive constant C_{MLMC} *such that for all* $\epsilon < e^{-1}$ *there are values* L *and* N_l *for which* $\hat{\mathcal{P}}_L^{(\mathrm{MLMC})}$ *satisfies*

$$\mathbb{E}\left[\left(\hat{\mathcal{P}}_L^{(\mathrm{MLMC})} - E\left[\mathcal{P}\right]\right)^2\right] < \epsilon^2,$$

with a computational complexity $\mathrm{Cost}^{(\mathrm{MLMC})}$ *with bound*

$$\mathrm{Cost}^{(\mathrm{MLMC})} \leq C_{MLMC} \cdot \epsilon^{-2}.$$

Condition (ii) is obviously true with our definition of $\hat{Y}_{\ell,N_\ell}$. Conditions (i) and (iv) are easy to satisfy and they hold for the two timestepping methods considered in the article. However, the quadratic variance decay in condition (iii) requires the correct coupling of paths between subsequent levels. The random variables ξ_n in Eqs. (8) and (9) that are used for calculating the paths on two coupled levels, have to be chosen correctly. For the symplectic Euler method in Eq. (8), this can be achieved by using a simple Brownian bridge, i.e. combining the random variables ξ_n and ξ_{n+1} on one level to obtain a random variable $\xi_n^{(\text{coarse})}$ on the next coarser level as

$$\xi_n^{(\text{coarse})} = \frac{1}{\sqrt{2}}(\xi_n + \xi_{n+1}). \tag{17}$$

As shown in [13], for the geometric Langevin method in Eq. (9), this is achieved by using

$$\xi_n^{(\text{coarse})} = \frac{e^{-\lambda(X_n)h}\xi_n + \xi_{n+1}}{\sqrt{e^{-2\lambda(X_n)h} + 1}}. \tag{18}$$

As is common in the MLMC literature, in the numerical section below, we carefully check condition (iii) by plotting the variance $\mathrm{Var}[Y_\ell] = N_\ell \cdot \mathrm{Var}[\hat{Y}_{\ell,N_\ell}]$ as a function of the step size h_ℓ.

4 MLMC with Reflective Boundary Conditions

As described in Section 2.1, we assume that particles are reflected elastically at the top and bottom of the domain. As shown in [16], this can be accounted for in the timestepping method by simply checking if a particle has left the domain after each timestep, and, if necessary, replacing $X_{n+1} \mapsto -X_{n+1}$ (or $X_{n+1} \mapsto 2H - X_{n+1}$ at the top boundary) and $U_{n+1} \mapsto -U_{n+1}$ in accordance with Eq. (6). This is shown schematically in Fig. 2a.

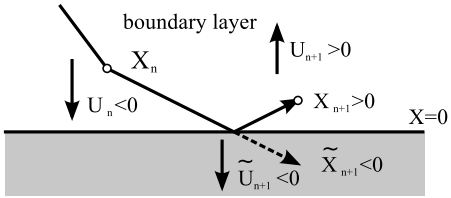

(a) Elastic reflection at the lower boundary

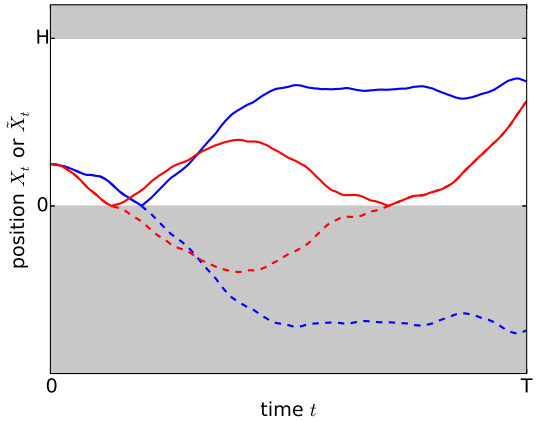

(b) Two paths in the original random variable X_t (solid lines) and the extended random variable $\tilde{X}_t$ (dashed lines)

Figure 2: Treatment of elastic boundary conditions

However, since the paths on subsequent levels are coupled in MLMC, care has to be taken to prevent the fine and coarse paths from diverging, which would result in violation of condition (iii) in the complexity theorem. This can be achieved by modifying Eqs. (17) and (18) appropriately. The key idea is illustrated by the following thought experiment, which we first explain for one reflective boundary (in Section 4.1.1 we show that the same treatment works for both boundaries by considering periodic extensions): imagine extending the domain of X to the whole real line and making it equivalent to the problem without boundary conditions. For this, an extended SDE in the variables $\tilde{X}$ and $\tilde{U}$ is formulated, which is related to the initial one by

$$\begin{pmatrix} U \\ X \end{pmatrix} = S(\tilde{X}) \begin{pmatrix} \tilde{U} \\ \tilde{X} \end{pmatrix}, \tag{19}$$

where $S(\tilde{X}_t) = (-1)^{n_{\text{refl}}(t)} = \operatorname{sign}(\tilde{X}_t)$ and $n_{\text{refl}}(t)$ is the number of reflections up to time t. Let $a(X,U)$ and $b(X)$ be some functions which are defined for positions $X \geq 0$ and any velocity U. For the SDE

$$dU_t = a(X_t, U_t)\,dt + b(X_t)\,dW_t, \qquad dX_t = U_t\,dt, \tag{20}$$

with reflection at $X = 0$, the equivalent extended SDE is

$$d\tilde{U}_t = A(\tilde{X}_t, \tilde{U}_t)\,dt + B(\tilde{X}_t)\,dW_t, \qquad d\tilde{X}_t = \tilde{U}_t\,dt. \tag{21}$$

The functions A and B in Eq. (21) are related to a and b in Eq. (20) by

$$A(\tilde{X}_t, \tilde{U}_t) = \begin{cases} a(\tilde{X}_t, \tilde{U}_t) & \text{if} \quad \tilde{X}_t \geq 0, \\ -a(-\tilde{X}_t, -\tilde{U}_t) & \text{if} \quad \tilde{X}_t < 0, \end{cases} \qquad B(\tilde{X}_t) = \begin{cases} b(\tilde{X}_t) & \text{if} \quad \tilde{X}_t \geq 0, \\ b(-\tilde{X}_t) & \text{if} \quad \tilde{X}_t < 0. \end{cases} \tag{22}$$

In Fig. 2b, both the solution $\tilde{X}_t$ of the extended SDE (21) and the solution X_t of the original SDE (20) are plotted. The solutions are related by Eq. (19).

One possible approach would now be to use the extended SDE (21), which does not have any boundary conditions, and then reconstruct the variables U_t, X_t by Eq. (19). However, it turns out that this is not necessary, and we can instead use the original set of equations if we carefully adapt the way the random variables are coupled in the MLMC algorithm.

4.1 Coupling between paths on different levels

To achieve this, we derive some relations between the extended and the original variables and the SDEs in Eqs. (20) and (21). Discretising Eq. (19) gives

$$\begin{pmatrix} U_n \\ X_n \end{pmatrix} = S_n \begin{pmatrix} \tilde{U}_n \\ \tilde{X}_n \end{pmatrix}, \tag{23}$$

where $S_n = (-1)^{n_{\text{refl}}(t_n)}$. By using this relation in Eq. (22), we obtain

$$A(\tilde{X}_n, \tilde{U}_n) = S_n a(S_n \tilde{X}_n, S_n \tilde{U}_n) = S_n a(X_n, U_n), \qquad B(\tilde{X}_n) = b(S_n \tilde{X}_n) = b(X_n). \tag{24}$$

Now consider the fine path on a MLMC level with timestep size h. Discretising the extended model in Eq. (21) with the symplectic Euler method gives

$$\tilde{U}_{n+1} = \tilde{U}_n + A(\tilde{X}_n, \tilde{U}_n)\, h + B(\tilde{X}_n)\sqrt{h}\, \tilde{\xi}_n, \qquad \tilde{X}_{n+1} = \tilde{X}_n + \tilde{U}_{n+1}\, h, \tag{25}$$

where $\tilde{\xi}_n \sim \text{Normal}(0,1)$ i.i.d. Replacing $(\tilde{U}_n, \tilde{X}_n)$ by (U_n, X_n) according to Eq. (23) results in an equation for the original variables U_n and X_n

$$\begin{aligned} S^{(f)}_{n+1} U_{n+1} &= S^{(f)}_n U_n + S^{(f)}_n a(X_n, U_n)\, h + b(X_n)\, \sqrt{h}\, \tilde{\xi}_n, \\ S^{(f)}_{n+1} X_{n+1} &= S^{(f)}_n X_n + S^{(f)}_{n+1} U_{n+1}\, h. \end{aligned} \tag{26}$$

To distinguish fine and coarse levels, we use the superscripts "(f)" and "(c)" on the reflection factors. Multiplying both sides by $\left(S^{(f)}_{n+1}\right)^{-1} = S^{(f)}_{n+1}$ and defining $s^{(f)}_{n+1} = S^{(f)}_{n+1} S^{(f)}_n$, the timestepping scheme can be expressed as

$$\begin{aligned} U_{n+1} &= s^{(f)}_{n+1} \left(U_n + a(X_n, U_n)\, h + b(X_n)\, \sqrt{h}\, S^{(f)}_n\, \tilde{\xi}_n \right), \\ X_{n+1} &= s^{(f)}_{n+1} \left(X_n + s^{(f)}_{n+1}\, U_{n+1}\, h \right). \end{aligned} \tag{27}$$

The same equation can be written down on the next coarser level, where we replace $h \mapsto 2h$, $S_n \mapsto S^{(c)}_n$ etc. Note that $s^{(f)}_{n+1}$ equals -1 if there is a reflection at time t_{n+1} or 1 otherwise. Hence the factors $s^{(f)}_{n+1}$ in Eq. (27) take care of the elastic reflection described above: if $X_{n+1} < 0$, velocity and position (U_{n+1}, X_{n+1}) get replaced by $(-U_{n+1}, -X_{n+1})$. In addition, the random forcing term proportional to $b(X_n)$ contains an additional factor $S^{(f)}_n$. This could be taken into account by sampling $\tilde{\xi}_n \sim \text{Normal}(0,1)$ and multiplying each sample by the correct factor $S^{(f)}_n$ as in Eq. (27). However, since the normal distribution is symmetric and $S^{(f)}_n = \pm 1$, on each individual level, the same result is obtained by directly using a random variable $\xi_n \sim \text{Normal}(0,1)$ and replacing $S^{(f)}_n \tilde{\xi}_n \mapsto \xi_n$ to obtain

$$U_{n+1} = s^{(f)}_{n+1} \left(U_n + a(X_n, U_n)\, h + b(X_n)\, \sqrt{h}\, \xi_n \right). \tag{28}$$

Formally the random variables ξ_n and $\tilde{\xi}_n$ are related by

$$\xi_n = S^{(f)}_n \tilde{\xi}_n \tag{29}$$

and agree in distribution since $S^{(f)}_n = \pm 1$. To couple the fine and coarse path correctly, it is necessary to require that the random variable $\tilde{\xi}^{(\text{coarse})}_n$ on the *coarse* paths fulfils the conditions in Eqs. (17) and (18). For the Euler scheme, this means that

$$\tilde{\xi}^{(\text{coarse})}_n = \frac{1}{\sqrt{2}} \left(\tilde{\xi}_n + \tilde{\xi}_{n+1} \right).$$

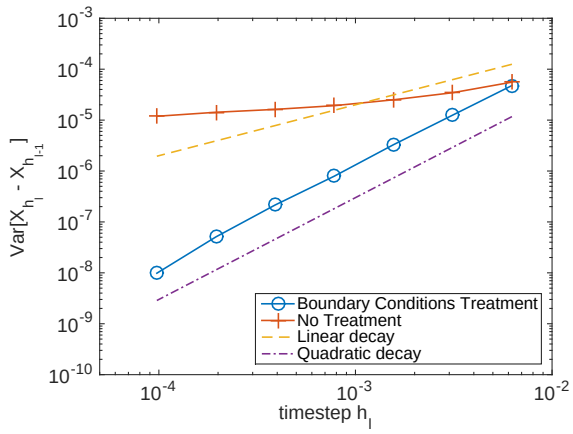

Figure 3: Geometric Langevin method - correct boundary condition treatment for mean particle position

Using Eq. (29) and the corresponding relation $\xi_n^{(\text{coarse})} = S_n^{(c)} \tilde{\xi}_n^{(\text{coarse})}$, this implies that reflection can be taken into account by replacing Eq. (17) by

$$\xi_n^{(\text{coarse})} = \frac{S_n^{(c)}}{\sqrt{2}} \left(S_n^{(f)} \xi_n + S_{n+1}^{(f)} \xi_{n+1} \right). \tag{30}$$

A similar argument shows that for the geometric Langevin method (18) needs to be replaced by

$$\xi_n^{(\text{coarse})} = S_n^{(c)} \frac{e^{-\lambda(X_n)h} S_n^{(f)} \xi_n + S_{n+1}^{(f)} \xi_{n+1}}{\sqrt{e^{-2\lambda(X_n)h} + 1}}. \tag{31}$$

Since this treatment of the reflective boundary conditions does not change the distribution of the normal random variables used in the coarse step, no extra bias is added and the telescoping sum (13) is preserved.

4.1.1 Reflection at the top boundary

To treat elastic reflection at both the top and bottom boundary, periodic extensions of the domain can be used. For $X \in \mathbb{R}$, let $n(X) = 2n$ if $X \in [(2n-1)H, (2n+1)H))$, $n \in \mathbb{Z}$ and $\eta(X) = X - n(X)H$ be a function that maps X to $[-H, H)$. Similarly to Eq. (22), the extended coefficients can be defined as

$$A(\tilde{X}_t, \tilde{U}_t) = \begin{cases} a(\eta(\tilde{X}_t), \tilde{U}_t) & \text{if} \quad \eta(\tilde{X}_t) \geq 0, \\ -a(-\eta(\tilde{X}_t), -\tilde{U}_t) & \text{if} \quad \eta(\tilde{X}_t) < 0, \end{cases} \quad B(\tilde{X}_t) = \begin{cases} b(\eta(\tilde{X}_t)) & \text{if} \quad \eta(\tilde{X}_t) \geq 0, \\ b(-\eta(\tilde{X}_t)) & \text{if} \quad \eta(\tilde{X}_t) < 0. \end{cases} \tag{32}$$

It is easy to verify that the function $\eta(X)$ reduces the problem to the case of reflecting only at $X = 0$ and the same analysis as in Section 4.1 applies. Therefore, no further changes are required in Eqs. (30) and (31) and $S_n^{(\cdot)}$ counts reflections at both boundaries.

Fig. 3 shows the variance $\text{Var}[\hat{Y}_{\ell, N_\ell}]$ of the difference process as a function of the timestep size. As can be seen from this plot, the quadratic decay required for condition (iii) in the complexity theorem is recovered. In the same figure, we also show the variance decay with naive (and incorrect) coupling of the paths at the reflective boundaries, i.e. when setting $S_n^{(\cdot)} = 1$. Due to the symmetry properties of the Gaussian increments, this still leads to the same unbiased estimators on each level, satisfying Assumptions (i), (ii) and (iv), but the variance of the difference decays with a rate that is not even linear and violates condition (iii).

Formally the discontinuity in the function $A(X, U)$ in the extended model might cause problems since the drift function is not continuous at $X = 0$. This problem could be fixed by using a smooth regularisation in

Section 2.1, which would ensure that higher derivatives of $\sigma_U(X)$ and $\tau(X)$ vanish at $X = 0$. However, as Fig. 3 demonstrates, it can be confirmed numerically that the method works for the model considered here even with the sharp regularisation in Eq. (7).

5 Smoothed indicator functions for concentration estimates

The concentration of particles which land in an interval $[a, b] \subset [0, H]$ can be expressed as the expected value of the indicator function defined in Section 2, i.e. $\mathbb{E}\left[\mathbb{1}_{[a,b]}(X_T)\right]$. However, since the indicator function is not continuous, this leads to problems in the MLMC method: it manifests itself in the fact that the variance on subsequent levels only decays linearly with the step size and condition (iii) in the complexity theorem is violated (this is consistent with the fact that the modified-equation analysis in [13], which is used to prove the complexity theorem there, also assumes that the quantity of interest is Lipschitz continuous). The numerical results below (Figs. 7b and 7a) confirm that this is indeed the case.

As we will demonstrate now, this issue can be fixed by approximating the indicator function by a suitable smooth function; any errors introduced by this approximation can be quantified and minimised systematically.

5.1 Smoothing polynomials

In order to recover the quadratic variance decay rates and satisfy condition (iii), we construct a smooth function that approximates the indicator function as described in [18]. By linearity, it is possible to write

$$\mathbb{E}\left[\mathbb{1}_{[a,b]}(X_T)\right] = \mathbb{E}\left[\theta(X_T - b)\right] - \mathbb{E}\left[\theta(X_T - a)\right] \qquad \text{with} \quad \theta(X) = \begin{cases} 1 & \text{for} \quad X \leq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{33}$$

In [18], to approximate the step function $\theta(X)$ the authors define

$$g_r(X) = \begin{cases} 1 & \text{if} \quad X < -1, \\ p_r(X) & \text{if} \quad -1 \leq X \leq 1, \\ 0 & \text{if} \quad X > 1. \end{cases} \tag{34}$$

Here, $p_r(X)$ is a polynomial of degree at most $r + 1$ such that $p_r(-1) = 1$, $p_r(1) = 0$ and the first r moments of $p_r(X)$ over the interval $[-1, +1]$ agree with the moments of the step function $\theta(X)$ on the same interval. This is guaranteed if

$$\int_{-1}^{1} s^j p_r(s)\, ds = \frac{(-1)^j}{j+1}, \qquad \text{for all} \;\; j = 0, \ldots, r-1.$$

We assume that r is chosen such that the density $\rho(X_T)$ of the particles at the final time is r times continuously differentiable on $[S_0 - \delta_0, S_1 + \delta_0]$ for some $S_0, S_1 \in \mathbb{R}$ and $\delta_0 > 0$ with $[a, b] \subset [S_0, S_1]$.

With g_r as defined in Eq. (34), it can be shown that

$$\sup_{s \in [S_0, S_1]} \left| \mathbb{E}\left[\theta(X_T - s)\right] - \mathbb{E}\left[g_r\left(\frac{X_T - s}{\delta} \right) \right] \right| \leq c \cdot \delta^{r+1}, \tag{35}$$

for some constant $c > 0$ and all $\delta \in (0, \delta_0]$ [18]. By choosing δ and r appropriately, it is possible to approximate any QoI based on the indicator function to arbitrary accuracy.

More specifically, $\mathbb{1}_{[a,b]}(X_T) = \theta(X_T - b) - \theta(X_T - a)$ is approximated by

$$P^{r,\delta}_{[a,b]}(X_T) \equiv g_r\left(\frac{X_T - b}{\delta} \right) - g_r\left(\frac{X_T - a}{\delta} \right). \tag{36}$$

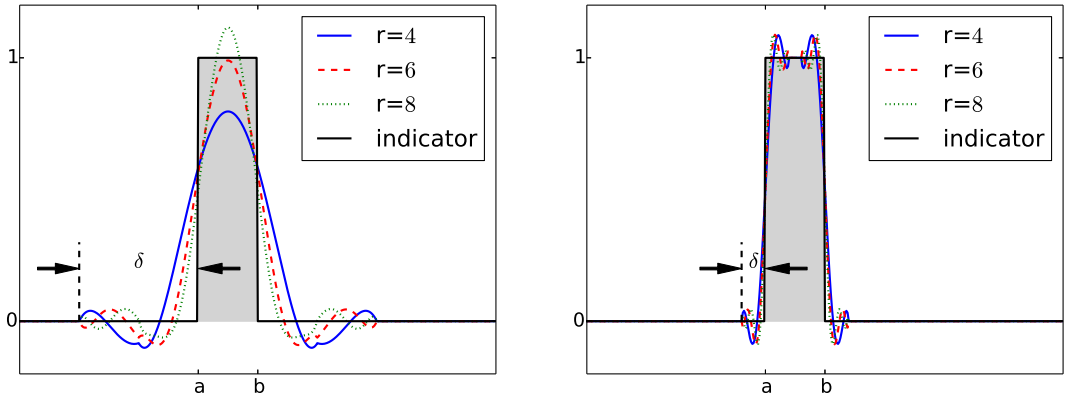

(a) $\delta = 0.1$ (b) $\delta = 0.02$

Figure 4: Smoothing polynomials $P^{r,\delta}_{[a,b]}(X)$ for different values of r, δ and for $a = 0.1$, $b = 0.15$. The indicator function is also shown.

This function is plotted for $\delta = 0.1$ and $\delta = 0.02$ and different polynomial degrees r in Figs. 4a and 4b.

Since (in contrast to the indicator function) $P^{r,\delta}_{[a,b]}$ is Lipschitz continuous, we expect to recover quadratic variance decay rates in MLMC. This is confirmed numerically in Section 6.2.1.

5.2 Error analysis

To quantify the error introduced by the smoothing process described in the previous section, we analyse the weak and strong errors of the problem with the smoothed indicator function. For ease of notation, we write $P(X) \equiv \mathbb{1}_{[a,b]}(X)$ and $P^{r,\delta}(X) \equiv P^{r,\delta}_{[a,b]}(X)$ and denote by $\hat{P}_L$ and $\hat{P}^{r,\delta}_L$ the Monte Carlo estimates of $\mathbb{E}[P]$ and $\mathbb{E}[P^{r,\delta}]$, which can be obtained either using StMC or MLMC estimators with stepsize $h_L = \frac{T}{M_L} = \frac{T}{M_0 2^L}$. After some straightforward algebra (see B) we obtain a bound on the bias (or weak) error

$$\left|\mathbb{E}[P] - \mathbb{E}[\hat{P}^{r,\delta}_L]\right| \leq 2c \cdot \delta^{r+1} + \left|\mathbb{E}[P^{r,\delta}] - \mathbb{E}[P^{r,\delta}_L]\right|. \tag{37}$$

In this equation, the second term describes the bias error introduced by the timestepping method. It is $\propto h_L$ for the first-order methods used in this work. The total mean-square error is

$$\begin{aligned}\mathbb{E}\left[\left(\hat{P}^{r,\delta}_L - \mathbb{E}[P]\right)^2\right] \leq \mathrm{Var}(\hat{P}^{r,\delta}_L) + \left(\mathbb{E}[P^{r,\delta}_L] - \mathbb{E}[P^{r,\delta}]\right)^2 + 4c^2 \cdot \delta^{2(r+1)} \\ + 4c \cdot \delta^{r+1}\left|\mathbb{E}[P^{r,\delta}_L] - \mathbb{E}[P^{r,\delta}]\right|,\end{aligned} \tag{38}$$

where c, δ and r are as described before and the quantity $\left|\mathbb{E}[P^{r,\delta}_L] - \mathbb{E}[P^{r,\delta}]\right| \propto h_L$ is the timestepping error.

From the strong-error bound (38), it can be observed that it suffices to let $\delta \to 0$ or $r \to \infty$ to ensure that the last two terms are dominated by the first two terms which represent the sampling error and the discretisation error as in the case of the indicator function without smoothing. Therefore, it is possible to systematically eliminate the smoothing error for practical calculations by choosing suitable r and δ. However, care has to be taken since the variance of Y_ℓ increases as the approximation to the indicator function improves. This can be observed numerically, for $r \to \infty$, in Fig. 7a. The behaviour of $\mathrm{Var}[Y_\ell]$ for $\delta \to 0$ is similar.

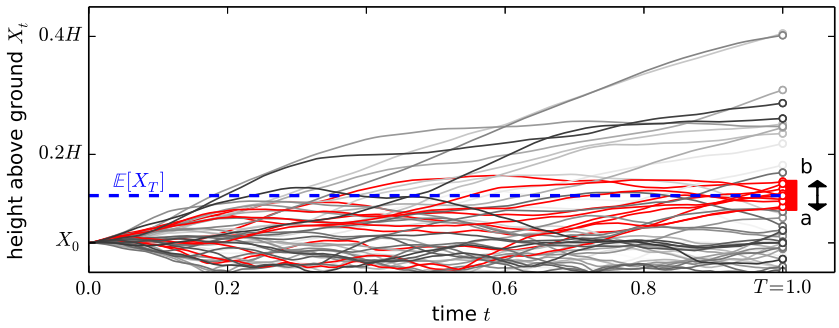

Figure 5: Sample trajectories from a typical model run. The expected value $\mathbb{E}[X_T]$ and the interval $[a, b]$ used for the concentration calculation are shown on the right.

6 Numerical Results

We now describe several numerical experiments to systematically quantify errors and to compare the efficiencies of the single and multilevel estimators with both timestepping methods. Results for three quantities of interest are presented: the particle position X_T at the final time T, the concentration in an interval $[a, b]$ and a piecewise constant approximation of the concentration field over the entire boundary layer. The latter is of particular interest in practical applications. In all cases, we first check the quadratic variance decay required in condition (iii) of the complexity theorem and then calculate the total computational cost as a function of the tolerance ϵ for the root-mean-square error. The number of timesteps M_0 for the coarsest level is chosen in advance (according to stability constraints), while the number of levels L is chosen so that the bias error is of size $\epsilon/\sqrt{2}$ (see [13] for details). To limit the total mean square error to ϵ^2, the second term in (16) is also reduced below $\epsilon^2/2$ by adaptively choosing the number of samples N_ℓ using on-the-fly estimators for $\mathrm{Var}[Y_\ell]$, see Eqn. (21) in [9].

All results were generated with a freely available object-oriented C++ code developed by the authors that can be downloaded under the GPL license from `https://bitbucket.org/em459/mlmclangevin`. The numerical results reported here were obtained with the version which has been archived as [26]. The code was compiled with version 4.8.4 of the GNU C++ compiler and run sequentially on an Intel i5-4460 CPU with a clock speed of 3.20GHz.

In the numerical experiments, we use $\kappa_\sigma = 1.3, \kappa_\tau = 0.5, u_* = 0.2\,\mathrm{m/s}$ and a boundary-layer depth of $H = 1\,\mathrm{km}$ in Eq. (5). All times, positions and velocities are expressed in units of the reference distance $X_{\text{ref}} = H = 10^3\,\mathrm{m}$, reference velocity $U_{\text{ref}} = 1\,\mathrm{m/s}$ and reference time $t_{\text{ref}} = X_{\text{ref}}/U_{\text{ref}} = 10^3 s \approx 17$ minutes. Particles are released at $X_0 = 0.05$ (corresponding to a height of $50\,\mathrm{m}$ above ground) with a velocity of $U_0 = 0.1$ (i.e., an upward release speed of $0.1\,\mathrm{m/s}$), but we also study the impact of the release height on performance in Section 6.3. Some sample trajectories for this set-up are shown in Fig. 5.

Unless explicitly stated otherwise, the regularisation height is $\epsilon_{\text{reg}} = 0.01$. In this case, $\tau(X)$ does not exceed a value of $20s$, which is consistent with the choice in the Met Office NAME model. The impact of ϵ_{reg} on the results is quantified in Section 7.1.

6.1 Particle Position

Since it is not affected by the additional complications due to smoothing of the indicator function, we first consider the expected value of the particle position at the final time, $\mathbb{E}[X_T]$. For both numerical methods, we use $M_0 = 40$ timesteps on the coarsest level and a final time $T = 1$ (corresponding to about 17 minutes).

The validity of condition (iii) in the complexity theorem is confirmed numerically in Fig. 6a where we plot the variance of $Y_\ell = X_{T,h_\ell} - X_{T,h_{\ell+1}}$ for decreasing timestep size h_ℓ. For both numerical methods, the

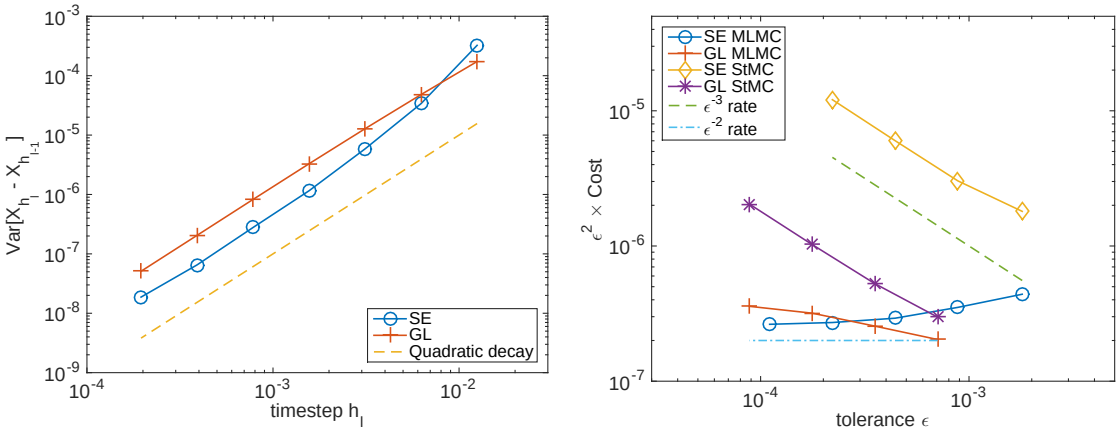

(a) Variance for different timestep sizes

(b) Computational cost for varying tolerance

Figure 6: Variance and total computational cost - Mean particle position

rate is quadratic, $\mathrm{Var}[Y_\ell] \propto h_\ell^2$. We observe that the Symplectic Euler (SE) method has a smaller variance than the Geometric Langevin method (GL) except for the largest timestep sizes. However, we also find that for a fixed value of h_ℓ, the bias error for the GL method is around $20\times$ smaller than for the SE algorithm.

In Fig. 6b, the total computational costs of the StMC and MLMC estimators for both discretisation methods are plotted as functions of the root-mean-square error tolerance ϵ. The StMC method shows an asymptotic scaling proportional to ϵ^{-3} for both discretisation schemes, as discussed in Section 3.2. The cost of the MLMC method is proportional to ϵ^{-2} in both cases, as predicted by the complexity theorem in Section 3.3. However, for a fixed ϵ the balance between bias and statistical error is different for the two discretisation methods: for fixed h_ℓ, the GL method has a smaller bias but larger variance error, so that, for a given ϵ, the SE method will require a smaller timestep. For the SE method, the maximal allowed timestep size is restricted by stability constraints, whereas for the GL algorithm there are no such restrictions.

Comparing only the MLMC results among themselves, both discretisation schemes perform similarly with only small differences. For the StMC algorithm, on the other hand, the GL method is clearly superior to the SE method by almost an order of magnitude across the range of tolerances ϵ considered. We conclude that while for larger tolerances ($> 10^{-3}$) the StMC method with GL timestepping is the most efficient, for small tolerances MLMC (with either discretisation scheme) leads to significant speed-ups.

6.2 Concentration

To estimate the concentration in a given interval $[a, b] \subset [0, 1]$, we calculate the expected value of the indicator function at the final time, i.e. $\mathbb{E}\left[\mathbb{1}_{[a,b]}(X_T)\right]$. The following results were obtained with an interval of width 0.05 centred at 0.1305 (i.e. $a = 0.1055$ and $b = 0.1555$; see Fig. 5). The centre of the interval roughly matches the expected value of X_T calculated in the previous section, $\mathbb{E}[X_T] = 0.1301 \pm 4 \cdot 10^{-4}$.

6.2.1 Sensitivity to smoothing parameters

As discussed in Section 5, to estimate concentrations with the MLMC algorithm it is useful to replace the discontinuous indicator function $\mathbb{1}_{[a,b]}$ by the smooth function $P^{r,\delta}_{[a,b]}$ defined in Eq. (36). In the following, we numerically quantify the errors introduced by this approximation and confirm the results from the theoretical analysis in Section 5.2. To verify condition (iii) in the complexity theorem and show that smoothing indeed leads to the correct variance decay, we plot in Figs. 7a and 7b the variances of Y_ℓ, for both timestepping methods, as functions of the timestep size h_ℓ. The quantities of interest are $P^{r,\delta}_{[a,b]}(X_T)$ with $\delta = 0.1$,

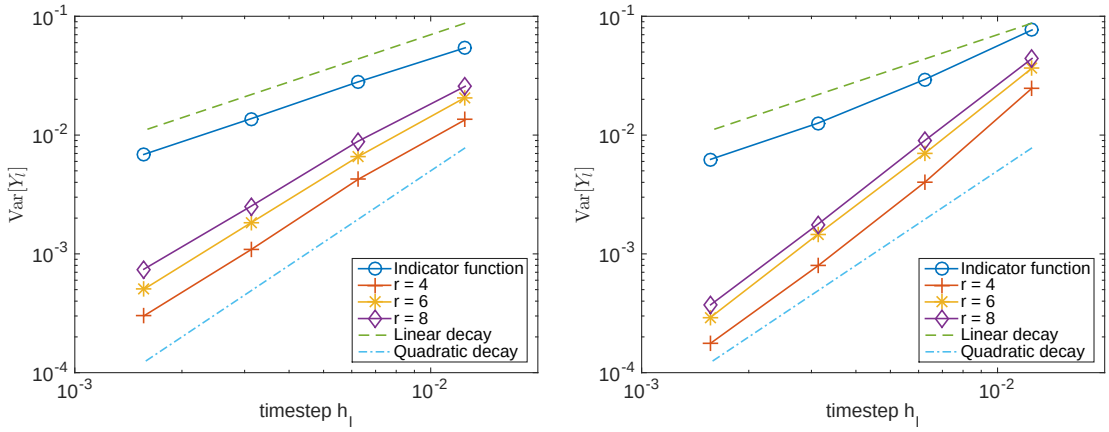

(a) Geometric Langevin

(b) Symplectic Euler

Figure 7: Variance - Concentration $P^{r,\delta}_{[a,b]}$ with $\delta = 0.1$

Quantity of interest	r	value		difference	
		GL	SE	GL	SE
$\mathbb{E}\left[P^{r,\delta}_{[a,b]}(X_T)\right]$	4	0.16596	0.16724	$7.1 \cdot 10^{-4}$	$8.0 \cdot 10^{-5}$
	6	0.16668	0.16794	$1.0 \cdot 10^{-5}$	$6.2 \cdot 10^{-4}$
	8	0.16682	0.16622	$1.5 \cdot 10^{-4}$	$1.1 \cdot 10^{-3}$
$\mathbb{E}\left[\mathbb{1}_{[a,b]}(X_T)\right]$		0.16667	0.16732		

Table 1: Estimates of the expected values for the indicator function and the smoothed function $P^{r,\delta}$ for different polynomial degrees r; in all cases $\delta = 0.1$ and $\epsilon = 10^{-3}$. The last two columns show the difference $\left|\mathbb{E}\left[P^{r,\delta}_{[a,b]}(X_T)\right] - \mathbb{E}\left[\mathbb{1}_{[a,b]}(X_T)\right]\right|$ between the smoothed functional and the indicator function.

for three different values of r, as well as the indicator function $\mathbb{1}_{[a,b]}$. These are the functions plotted in Fig. 4a. For both plots the tolerance on the root-mean-square is fixed to $\epsilon = 10^{-3}$. We observe quadratic dependence $\mathrm{Var}[Y_\ell] \propto h_\ell^2$ for the smoothed QoIs, while for the non-smooth indicator function the variance depends linearly on h_ℓ. For the GL method, the rates seem to be even superquadratic for larger timestep sizes. The absolute value of the variance grows as r increases and this is intuitively expected since for larger r the function $P^{r,\delta}_{[a,b]}$ approximates the non-smooth indicator function better.

To show that the error introduced by replacing the indicator function by a smooth approximation is indeed small, we tabulate estimates of the expected values of the indicator function and of $P^{r,\delta}_{[a,b]}$, for different values of r, in Tab. 1. As can be seen from the last two columns of this table, the difference is less than $1.1 \cdot 10^{-3}$ in all cases. Note that we are only requiring the root mean square error of the estimator $\hat{P}_{[a,b]}$ to be less than $\epsilon = 10^{-3}$. A particular instance of the estimator may therefore have a larger error, so that we can conclude that the smoothing error is negligible for the choices of δ and r in Tab. 1.

6.2.2 Cost comparison between Standard and Multilevel Monte Carlo

Having confirmed that the error from approximating the indicator function by a smooth polynomial can be made sufficiently small, we now compare the cost of the StMC and MLMC methods as a function of the tolerance on the root-mean-square error. Based on the results in the previous section, we fix $\delta = 0.1$ and $r = 4$. For both discretisation methods, the number of timesteps on the coarsest level is $M_0 = 80$.

In Fig. 8, the total computational cost for the StMC and MLMC method is plotted for different tolerances

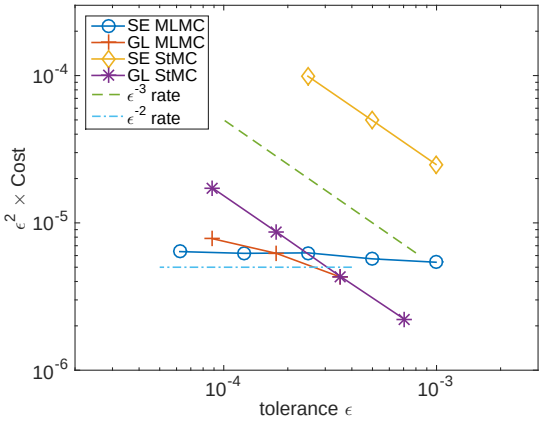

Figure 8: Computational cost for the concentration (using $P^{r,\delta}_{[a,b]}(X_T)$ with $r = 4$ and $\delta = 0.1$).

on the root-mean-square error ϵ. We see a similar picture as in Section 6.1. While for larger tolerances, the StMC with GL is cheaper than the MLMC algorithms, for tighter tolerances the MLMC algorithm becomes superior. We observe the expected growth $\propto \epsilon^{-3}$ for the two StMC methods. Again, comparing only the StMC algorithms, for a given ϵ, there is a large difference in the efficiency between the two discretisation methods, with the GL method being more than an order of magnitude faster. This is due to a large difference in the bias error, which is around $20\times$ smaller for a given timestep size h_ℓ. For the MLMC methods, there is again little dependence on the discretisation method. In both cases, the cost for MLMC grows at the expected rate of ϵ^{-2}, at least asymptotically. The slightly faster growth for the GL MLMC for large tolerances is probably due to the small number of levels (only 2 for the rightmost data point).

6.3 Probability Density Function

Since the main output of operational dispersion models like NAME is a box-averaged concentration field, we finally calculate an approximation to the probability density function (p.d.f.) at the final timestep. For this a set of equidistant points $0 = a_0 < a_1 < a_2 < \cdots < a_{k-1} < a_k = H$ is chosen and the particle concentration is estimated for each interval (or bin) $[a_i, a_{i+1}]$. In the Monte Carlo algorithm, this is achieved by considering the following vector-valued quantity of interest:

$$\boldsymbol{\mathcal{P}} = (\mathcal{P}_0, \mathcal{P}_1, \ldots, \mathcal{P}_k) = \left(P^{r,\delta}_{[a_0,a_1]}(X_T), P^{r,\delta}_{[a_1,a_2]}(X_T), \ldots, P^{r,\delta}_{[a_{k-1},a_k]}(X_T) \right), \tag{39}$$

which is related to the p.d.f. or particle concentration $\rho(X)$ by

$$\int_a^b \rho(X)\, dX = \mathbb{E}[\mathbb{1}_{a,b}] \approx \mathbb{E}[P^{r,\delta}_{[a,b]}] \qquad \text{for all intervals } [a,b] \subset [0,H].$$

As above, indicator functions are replaced by smoothed polynomials $P^{r,\delta}_{[a,b]}$ to guarantee quadratic variance decay in the MLMC algorithm. The QoI is now vector-valued, but a scalar quantity is required to estimate the optimal number of samples N_ℓ on each level of the MLMC algorithm (see [9, 10] for details); we choose the maximal absolute variance over all intervals for this scalar. More specifically, let $\mathcal{P}_{\ell,i} = P^{r,\delta}_{[a_i,a_{i+1}]}(X_{T,\ell})$ and $Y_{\ell,i} = \mathcal{P}_{\ell,i} - \mathcal{P}_{\ell-1,i}$, for $\ell = 0, \ldots, L$ and $i = 0, \ldots, k$. Then, we use $\max_{i=0,k} |\mathrm{Var}[Y_{\ell,i}]|$ in place of $\mathrm{Var}[Y_\ell]$ to estimate N_ℓ on level ℓ .

Fig. 9 shows the binned probability density function of X_T obtained with the SE method for a box size of 0.05. The leftmost plot is obtained with $X_0 = 0.05$, corresponding to a release height of $50\,\mathrm{m}$ in physical units; this is the same initial condition which was used to obtain all previous results in Sections 6.1 and 6.2.

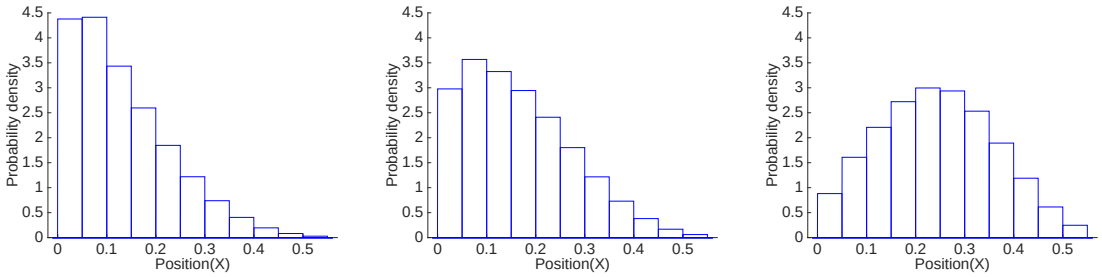

Figure 9: Binned concentration field for release height 0.05 (left), 0.10 (center) and 0.20 (right). The results were obtained with the MLMC method and symplectic Euler timestepping.

release height	L	h_L	StMC SE	StMC GL	StMC speed-up	MLMC SE	MLMC GL	Speed-up MLMC-StMC
0.05 ($50\,\mathrm{m}$)	7	$9.8 \cdot 10^{-5}$	3846.55	239.79	$16.0\times$	193.53	198.50	$1.2\times$
0.10 ($100\,\mathrm{m}$)	5	$3.9 \cdot 10^{-4}$	1282.66	201.11	$6.4\times$	109.01	128.10	$1.8\times$
0.20 ($200\,\mathrm{m}$)	3	$1.6 \cdot 10^{-3}$	200.73	142.61	$1.4\times$	51.10	67.99	$2.8\times$

Table 2: Computation cost for different release heights. Results are shown both for the geometric Langevin (GL) and for the symplectic Euler method (SE). All times are given in seconds and the last column shows the speed-up of the fastest MLMC method relative to the fastest StMC method. We also give the number of levels L and the fine level timestep size h_L to achieve a tolerance of $\epsilon = 4 \cdot 10^{-3}$ for each release height.

However, since we found that the computational cost is very sensitive to the release height, we also report results for $X_0 = 0.1$ and $X_0 = 0.2$, corresponding to $100\,\mathrm{m}$ and $200\,\mathrm{m}$ above ground, respectively.

We require that the root-mean-square error for the particle concentration $\rho(X)$ does not exceed a tolerance of $\epsilon = 4 \cdot 10^{-3}$. As before the number of timesteps on the coarsest level is $M_0 = 80$. We observe that the bias error decreases with increasing release height, since the particles spend less time near the ground where $\tau(X)$ varies strongly. Hence as X_0 increases, the size of the timestep h_L that is required to reduce the bias error below the tolerance $\epsilon/\sqrt{2}$ on the finest level also increases. The timesteps sizes h_L used in Fig. 9 using MLMC with the SE scheme are shown in Tab. 2 together with the computational cost that is required to reduce the root-mean-square error below a tolerance of $\epsilon = 4 \cdot 10^{-3}$ with all the different estimators.

Looking at the StMC algorithm, we note again that the performance is improved dramatically by using the GL method instead of SE. This improvement is particularly dramatic for small release heights where the GL algorithm is about $16\times$ faster than the SE method, which agrees with the speed-ups observed for other quantities of interest in the previous sections. The difference is less pronounced for higher release heights, but even for $X_0 = 0.2$ GL timestepping reduces the solution time by around 40%. The reason for this is the reduction in bias error due to better treatment of the large and strongly varying profile $\lambda(X)$ by the GL method which enables us to use a larger timestep.

For the MLMC algorithm, the performance depends less on the timestepping method, with the SE giving slightly better results. Similar to the StMC algorithm, the total runtime is reduced as the release height increases. The relative performance gains of MLMC over StMC also depend on X_0. While for the lowest release height, the MLMC-SE method is only $1.2\times$ faster than StMC-GL, the speed-up is $2.8\times$ for $X_0 = 0.2$.

In summary, compared to the reference StMC-SE algorithm used in many Lagrangian atmospheric dispersion models, both the new GL timestepping method and the MLMC algorithm reduce the runtime significantly when computing concentration fields.

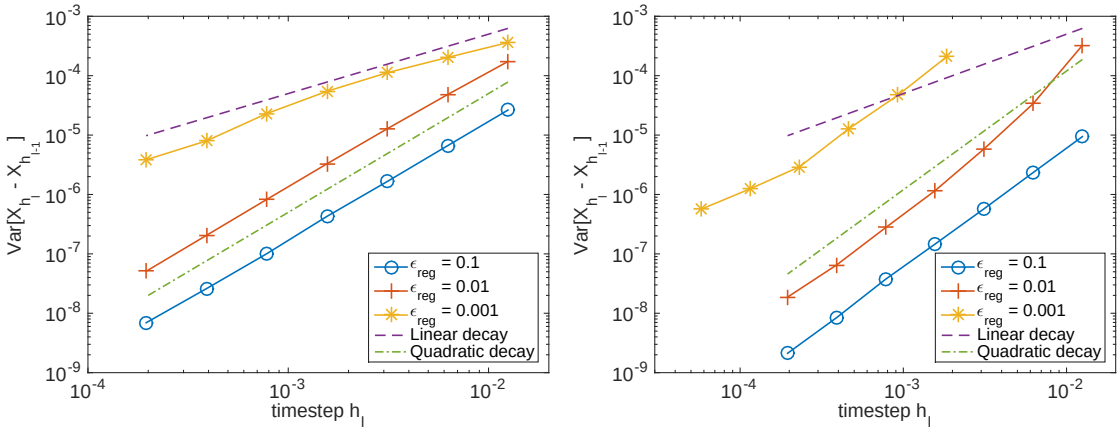

(a) Geometric Langevin

(b) Symplectic Euler

Figure 10: Variance rates - Mean particle position

7 Adaptive timestepping

Since the velocity decorrelation time $\tau(X)$ varies strongly over the boundary layer and limits the timestep size near the ground, we also investigate potential improvements from adaptive timestepping methods. By using small timesteps only in those parts of the domain where $\tau(X)$ is large, those methods can potentially reduce the overall runtime. For StMC adaptive timestepping is easy to implement and is currently used in the NAME dispersion model. Adaptive timestepping methods can also be used in the MLMC algorithm. We briefly review the technique proposed in [25] and adapt it for the timestepping methods used here.

7.1 Sensitivity to regularisation height

As the largest possible value of $\tau(X)$ depends on the regularisation height ϵ_{reg} introduced in Section 2.1, we first investigate the impact of this parameter on our results. As before, the first thing to check is the quadratic variance decay $\text{Var}[Y_\ell] \propto h_\ell^2$ required for condition (iii) of the complexity theorem. In Figs. 10a and 10b, the variance is shown as a function of timestep size h_ℓ for both timestepping methods as ϵ_{reg} is varied. In both cases, the quantity of interest is the final particle position X_T and the same set-up as in Section 6.1 is used. For GL, the number of timesteps on the coarsest level is $M_0 = 40$ and up to $L = 7$ levels are used. For SE, the timestep sizes are limited by stability constraints, i.e. $h < C/\max_{X\in[0,H]}\{\lambda(X)\} = C/\lambda(\epsilon_{\text{reg}})$ for some constant C. This restriction becomes more severe as the size of the regularisation parameter decreases. For $\epsilon_{\text{reg}} = 0.1$ and $\epsilon_{\text{reg}} = 0.01$, we use the same set-up as for the GL methods, but for $\epsilon_{\text{reg}} = 0.001$ a larger number of $M_0 = 270$ on the coarsest level and a reduced number of $L = 6$ is used to ensure that the method is stable on all levels.

From those plots, it can be observed that the variance decay is indeed quadratic for $\epsilon_{\text{reg}} \geq 0.01$ for the values of h_ℓ considered, but $\text{Var}[Y_\ell] \propto h_\ell$ for the smallest ϵ_{reg}. This deterioration in the variance decay rate can be explained using the theory of modified equations [27, 28, 13] in A: since $\lambda \to \infty$ for $X \to 0$, the variance decay rate is quadratic only if $h_\ell \cdot \epsilon_{\text{reg}}^{-2}$ is not too large. This is because for larger h_ℓ such that $h_\ell \cdot \epsilon_{\text{reg}}^{-2} \gtrsim 1$, heuristically we can only hope for linear variance decay with $\text{Var}[Y_\ell] \propto h_\ell$ since the size of the term $h_\ell v_1^{(U)}$ in Eq. (44) is $\mathcal{O}(1)$ and not $\mathcal{O}(h_\ell)$ and therefore the modified equation expansion breaks down. However, asymptotically, as h_ℓ goes to 0 (for fixed ϵ_{reg}), we recover the expected quadratic variance decay rate $\text{Var}[Y_\ell] \propto h_\ell^2$ as soon as $h_\ell \cdot \epsilon_{\text{reg}}^{-2} \ll 1$. We stress, however, that for the regularisation height used in all experiments above, which roughly agrees with what is used in the Met Office NAME model, the variance indeed behaves like $\text{Var}[Y_\ell] \propto h_\ell^2$ for practically relevant values of h_ℓ and hence condition (iii) in

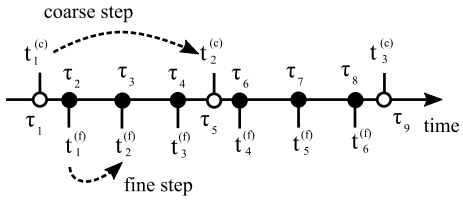

Figure 11: Non-nested timesteps in the adaptive MLMC algorithm.

the complexity theorem is satisfied.

For both discretisation methods the main reason for the increase in the variance is the large value of λ relative to the timestep size when the particle is close to the lower boundary. By decreasing the timestep size, the variance decreases but at the same time the total cost increases. We therefore now look at how this effect can be mitigated with adaptive timestepping.

7.2 Adaptive timestepping

Various methods for adaptive timestepping in the MLMC method have been proposed in the literature [25, 34, 35, 36]. Instead of using a fixed timestep h for all timesteps, in an adaptive algorithm the timestep size is adjusted at every time t_n. Here we use

$$h_n^{(\text{adaptive})} := \min\left\{h, \frac{\lambda(X_{\text{adapt}})}{\lambda(X_n)}h\right\}, \tag{40}$$

where X_n is the current particle position and X_{adapt} is a reference height. Since $\lambda(X)$ decreases with height, this choice guarantees that the product of the adaptive timestep size and $\lambda(X_n)$ never exceeds $h\,\lambda(X_{\text{adapt}})$, which can then be chosen such that the SE method is stable. Since the telescoping sum in Eq. (13) is preserved, this choice of adaptive timestep is independent of the MLMC level and therefore does not introduce any additional bias [25]. Since $\lambda(X)$ is a decreasing function of X, it is only necessary to adapt the timestep size for particles below the reference height X_{adapt}. However, since $\lambda \to \infty$ close to the ground, it is still necessary to regularise $\lambda(X)$ as described in Section 2.1. Without regularisation, the timestep size can become arbitrarily small, and it can take an infinite amount of time to calculate even a single trajectory.

When using adaptivity within MLMC, the coarse and fine timesteps are not necessarily nested. However, it is easy to adapt the MLMC algorithm to account for this [25]. For this, the time interval $[0, T]$ is divided into a number of intervals $[\tau_j, \tau_{j+1}]$ such that $\tau_0 = 0$, $\tau_N = T$ and each of the points τ_j is either a fine or a coarse time point (see Fig. 11, which has been adapted from Fig. 1 in [25]). More explicitly, define

$$\tau_j = \begin{cases} t_i^{(c)} & \text{for coarse time points,} \\ t_{i'}^{(f)} & \text{for fine time points.} \end{cases} \tag{41}$$

A normal random variable $\xi_j \sim \text{Normal}(0,1)$ is associated with each interval $[\tau_j, \tau_{j+1}]$. Then, for each coarse or fine time interval $[t_i^{(x)}, t_{i+1}^{(x)}]$ (where $x = c, f$), those random variables can be combined to construct a random increment $\Delta W(t_i^{(x)}, t_{i+1}^{(x)}) \sim \text{Normal}(0, t_{i+1}^{(x)} - t_i^{(x)})$ with the correct distribution. For SE with reflection, the expression for the coarse path is

$$\Delta W(t_i^{(c)}, t_{i+1}^{(c)}) = S_i^{(c)} \sum_{j=j_-}^{j_+-1} S_j^{(f)} \xi_j \sqrt{\tau_{j+1} - \tau_j}, \quad \text{where } \tau_{j_-} = t_i^{(c)} \text{ and } \tau_{j_+} = t_{i+1}^{(c)}. \tag{42}$$

Here, the reflection factor is $S_j^{(\cdot)} = (-1)^{n_{\text{refl}}(\tau_j)}$, where, as before, $n_{\text{refl}}(\tau_j)$ counts the number of reflections up to time τ_j. The corresponding expression $\Delta W(t_i^{(f)}, t_{i+1}^{(f)})$ for the fine path is the same as in (42), with

the only difference that the reflection factors $S_i^{(c)}$ and $S_j^{(f)}$ do not appear in this case. Since $\Delta W(t_i^{(x)}, t_{i+1}^{(x)})$ is a sum of Gaussian random variables, it is easy to see that this variable has a variance of $t_{i+1}^{(x)} - t_i^{(x)}$ and can be used instead of $\sqrt{h}\,\xi_n$ for the velocity increment in (8). The crucial observation is that if a fine interval overlaps with a coarse interval, they share some of the random increments and this ensures the correct coupling between fine and coarse paths. The expression in (42) reduces to the standard expression for the case that the coarse points coincide with fine points. For GL, the expression in (42) has to be modified to

$$\Delta W(t_i^{(c)}, t_{i+1}^{(c)}) = S_i^{(c)} \sum_{j=j_-}^{j_+-1} \left\{ S_j^{(f)} \xi_j \sqrt{\frac{1-\exp[-2\lambda(X(\tau_j))\cdot(\tau_{j+1}-\tau_j)]}{2\lambda(X(\tau_j))}} \right. \\ \left. \cdot \exp\left[-\sum_{k=j+1}^{j_+-1} \lambda(X(\tau_k))\cdot(\tau_{k+1}-\tau_k) \right] \right\}. \tag{43}$$

This increment replaces $\sqrt{\frac{1-\exp[-2\lambda(X_n)h]}{2\lambda(X_n)}}\xi_n$ in (9). Again the corresponding expression for the fine path is obtained from Eq. (43) by removing the reflection factors $S_i^{(c)}$ and $S_j^{(f)}$. It can be verified that this is equivalent to (18) in the case that the coarse timesteps coincide with fine timesteps.

7.3 Numerical results

We now present some numerical results to quantify potential improvements from adaptive timestepping. The adaptive MLMC algorithm described in the previous section is used to calculate the expected value of the particle position X_T at the final time with the same set-up as in Section 6.1 and $X_{\text{adapt}} = 0.05$. As before, we first confirm that condition (iii) in the complexity theorem is satisfied. The variance $\text{Var}[Y_\ell]$ as a function of the timestep size h_ℓ is shown for both SE and GL integrators in Figs. 12a and 12b. As expected, both for the uniform and adaptive timestepping methods, the variance decays quadratically with the (maximal) timestep size, $\text{Var}[Y_\ell] \propto h_\ell^2$. In both cases, adaptive timestepping reduces the variance for a given h_ℓ.

The total computational cost for both methods is compared in Figs. 13a and 13b. While the results for SE are as expected and adaptive timestepping reduces the overall cost, the opposite can be observed for GL; here uniform timestepping reduces the runtime (albeit not by much as $\epsilon \to 0$). Looking at the bias error for a given (maximal) timestep size, we find that for the SE method this bias is about $10\times$ smaller with adaptive timestepping. This and the fact that the variance is reduced (see Fig. 12a) explains the results in Fig. 13a. For the GL method, on the other hand, adaptive timestepping *increases* the bias by about $3\times$. This is a counter-intuitive result, since all individual timesteps in the adaptive algorithm are smaller than for the corresponding uniform implementation. For a (deterministic) ODE, it would then be easy to prove that adaptivity reduces the discretisation error. For the SDE considered here, however, the theoretical analysis is significantly harder and beyond the scope of this work. To exclude any bugs in the code, we have carried out several tests. For example, we have confirmed that for the corresponding ODE equations ($\sigma_U = 0$) adaptivity reduces the bias error also for the GL method. The phenomenon appears to be linked to the divergence of $\lambda(X)$ close to the ground and might require careful tuning of both X_{adapt} and ϵ_{reg} to see any benefits from adaptive timestepping.

8 Conclusions and future work

In this paper, we presented two improvements to Monte Carlo algorithms for the solution of SDEs encountered in atmospheric dispersion modelling: the geometric Langevin (GL) method uses a splitting approach to avoid numerical instabilities due to very small velocity decorrelation times close to the ground. Multilevel Monte Carlo reduces the algorithmic complexity as a function of the tolerance ϵ on the root-mean-square error from ϵ^{-3} to ϵ^{-2}. Both methods were applied to the simulation of the vertical spread and diffusion

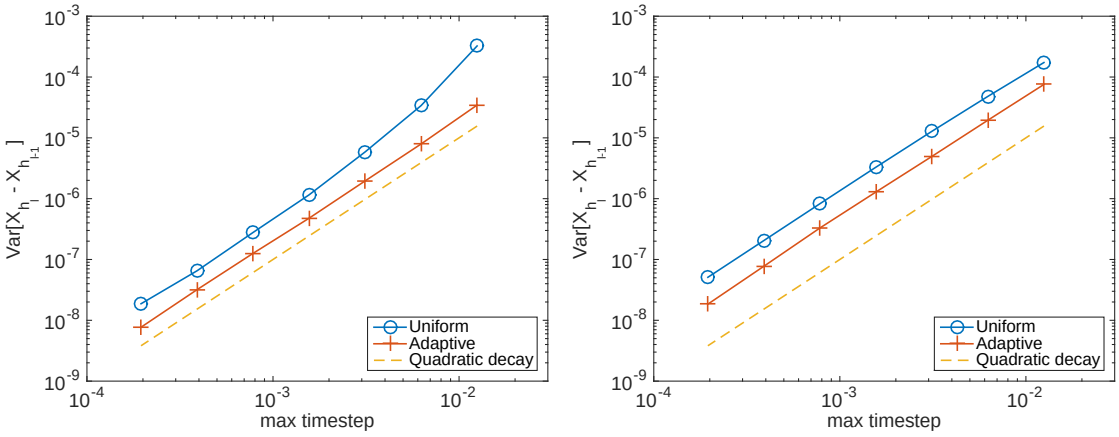

(a) Symplectic Euler (b) Geometric Langevin

Figure 12: Variance $\mathrm{Var}[Y_\ell]$ as function of the timestep size

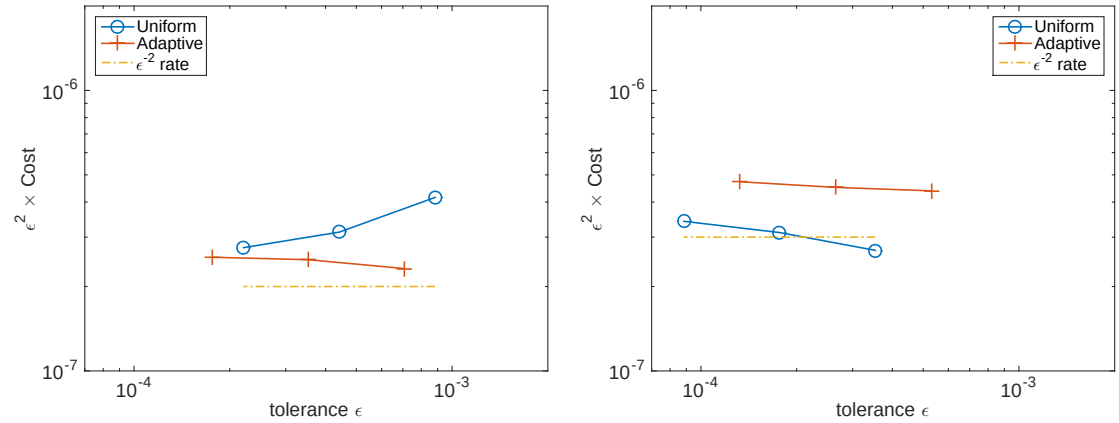

(a) Symplectic Euler (b) Geometric Langevin

Figure 13: Computational cost of adaptive- and non-adaptive timestepping algorithms for different tolerances ϵ on the root-mean-square error.

of particles in a turbulent boundary layer. The intricacies of the application considered here requires careful adaptation of the algorithms to account for suitable boundary conditions at the top and bottom of the atmosphere and to treat the divergence of the inverse velocity decorrelation time near the ground. When predicting particle concentrations, a smoothed indicator function has to be used and we confirmed that the additional errors introduced by the smoothing are small.

If the initial condition is chosen such that particles are released at a relatively low height, the GL algorithm reduces the computational cost (for the same accuracy) by around at least an order of magnitude for all quantities of interest considered in our numerical experiments (as compared to Symplectic Euler). The main reason for this is a significant reduction in the bias error since the GL algorithm is able to deal with large values of the inverse velocity decorrelation time $\lambda(X)$. For small tolerances ϵ, the MLMC algorithm leads to further speed-ups, which are particularly pronounced if the particles are released at a larger height. When calculating piecewise constant approximations to the particle concentration field, the MLMC algorithm is nearly three times faster than the standard Monte Carlo algorithm even when used in combination with the improved GL timestepping method.

In a preliminary study we also explored the use of adaptive timestepping and find that this can lead to a further, but less dramatic, reduction in runtime for the symplectic Euler method. The unexpected results observed for the GL algorithm require further investigation and careful tuning of parameters.

Of central importance for the application of the MLMC algorithm is the correct dependence of the multilevel variance on the timestep size. Numerically we observed the correct behaviour, but a theoretical proof is more challenging due to the non-smooth regularisation of the profile functions. Although this problem can be addressed by using an equivalent smooth regularisation instead, a proof relies on further technical details such as boundedness of all moments of U_t, which would then allow the application of the modified equation analysis in [13]. The treatment of reflective boundary conditions introduces discontinuities, which do not have any adverse effects on our numerical results but require a careful treatment in a rigorous mathematical analysis. Since the focus of this work is the introduction of new numerical methods to atmospheric dispersion modelling, this would be beyond the scope of this article.

In addition to investigating those mathematical details more rigorously, there are various ways of extending our work. One obvious avenue to explore is the use of quasi Monte Carlo methods (QMCs), which potentially allow a reduction of the computational complexity beyond the standard Monte Carlo limit of ϵ^{-2}, which is set by the central limit theorem (see e.g. [37]). By using a randomly shifted point lattice, any additional bias from the QMC sampling can be removed.

So far we have treated turbulence in all space directions independently and focussed on motion in the vertical direction where the atmospheric conditions vary strongly. Although similar approximations are made in operational models, ultimately we would like to apply our approach to a full three-dimensional scenario, for example to study the global spread of ash clouds, which would require the use of a more complicated, 3d turbulence model described in [29]. One particular question to be addressed is as to whether the separation of particle pairs by the resolved flow field will have any impact on the efficiency of MLMC.

It is our hope that eventually operational models such as NAME will make use of the Multilevel Monte Carlo method. NAME models a plethora of additional physical effects such as gravitational settling, chemical interactions, radioactive decay chains and gamma radiation from non-local radioactivity, which are computationally expensive. While in this paper the multilevel approach was described in terms of a hierarchy of levels of increasing timestep size, the first coarsening step might consist of using the same timestep size but not including all physical processes. Conceptually this is somewhat similar to the difference between p- and h- refinement in finite element methods.

Acknowledgements

The PhD of Grigoris Katsiolides is supported by an EPSRC CASE studentship, which is partially funded by the Met Office.

A Modified equations and regularisation

Our theoretical tools for proving the complexity theorem rely on the theory of modified equations [27, 28] and their application to the MLMC method as described in [13]; this approach is an alternative to the usual strong convergence proofs. We discuss this analysis here since it can be used to justify the use of a regularisation cut-off at the upper and lower boundary as outlined in Section 4.

The key idea is as follows: under suitable conditions, the numerical solution X_n of the initial SDE (4) is the weak second-order approximation of the solution of a set of modified equations – recall that X_n converges only with order one to the original equations. If those modified equations exist, it can be shown that condition (iii) in the complexity theorem holds [13].

For our model, the modified equations have the form

$$\begin{aligned} dU &= \left(-\lambda(X)U - \frac{\partial V(X,U)}{\partial X} + h v_1^{(U)}(X,U)\right) dt + \left(\sigma(X) + h\Sigma_1^{(U)}(X,U)\right) dW_t, \\ dX &= \left(U + h v_1^{(X)}(X,U)\right) dt + h\Sigma_1^{(X)}(X,U) dW_t. \end{aligned} \tag{44}$$

The functions $v_1^{(U)}$, $v_1^{(X)}$, $\Sigma_1^{(U)}$ and $\Sigma_1^{(X)}$ depend on the numerical timestepping method and can be expressed explicitly in terms of $\tau(X)$ and $\sigma_U(X)$. However, these modified equations are only a first-order approximation of Eq. (4) if the $\mathcal{O}(h)$ terms in (44) are indeed small. Unfortunately, this is not the case near the boundary of the domain. In particular, by expanding

$$v_1^{(U)}(X,U) = X^{-2} \cdot f^{(v,U)}(U) + \ldots, \qquad v_1^{(X)}(X,U) = X^{-1} \cdot f^{(v,X)}(U) + \ldots, \tag{45}$$

$$\Sigma_1^{(U)}(X,U) = X^{-3/2} \cdot f^{(\Sigma,U)}(U) + \ldots, \qquad \Sigma_1^{(X)}(X,U) = X^{-1/2} \cdot f^{(\Sigma,X)}(U) + \ldots \tag{46}$$

it can be shown that these terms diverge at the lower boundary ($X \to 0$), for both the symplectic Euler and geometric Langevin. Similar expressions can be derived at the upper boundary ($X \to H$), but for simplicity we do not write them down here.

However, if we regularise the profiles $\tau(X)$ and $\sigma_U(X)$ as described in Section 2.1, then the divergent terms $X^{-\alpha}$ are replaced by the constant $\epsilon_{\text{reg}}^{-\alpha}$. Thus, the terms $v_1^{(U)}$, $v_1^{(X)}$, $\Sigma_1^{(U)}$ and $\Sigma_1^{(X)}$ can all be made $\mathcal{O}(1)$ by a suitable choice of ϵ_{reg}. The existence of the modified equations then implies the quadratic variance decay required by condition (iii) in the complexity theorem, under the condition that U has bounded moments. This can be seen in Fig. 10a: unless ϵ_{reg} is chosen too small, the variance decay is indeed quadratic.

B Derivation of error bounds for smoothed indicator function

To derive inequality (37), we first consider the following difference between the expected values of the indicator function P and the smoothed out function $P^{r,\delta}$:

$$\left|\mathbb{E}[P] - \mathbb{E}[P^{r,\delta}]\right| \leq \left|\mathbb{E}\left[\theta(X_T - b)\right] - \mathbb{E}\left[g_r\left(\tfrac{X_T - b}{\delta}\right)\right]\right| + \left|\mathbb{E}\left[\theta(X_T - a)\right] - \mathbb{E}\left[g_r\left(\tfrac{X_T - a}{\delta}\right)\right]\right| \leq 2c\delta^{r+1}. \tag{47}$$

Using this expression and $P - \hat{P}_L^{r,\delta} = (P - P^{r,\delta}) + (P^{r,\delta} - \hat{P}_L^{r,\delta})$, the bound on the bias (or weak error) in (37) follows from the triangle inequality

$$\left|\mathbb{E}[P] - \mathbb{E}[\hat{P}_L^{r,\delta}]\right| \leq \left|\mathbb{E}[P] - \mathbb{E}[P^{r,\delta}]\right| + \left|\mathbb{E}[P^{r,\delta}] - \mathbb{E}[\hat{P}_L^{r,\delta}]\right| \leq 2c\delta^{r+1} + \left|\mathbb{E}[P^{r,\delta}] - \mathbb{E}[P_L^{r,\delta}]\right|.$$

To prove (38), we write

$$\hat{P}_L^{r,\delta} - \mathbb{E}[P] = \left(\hat{P}_L^{r,\delta} - \mathbb{E}[P_L^{r,\delta}]\right) + \left(\mathbb{E}[P_L^{r,\delta}] - \mathbb{E}[P^{r,\delta}]\right) + \left(\mathbb{E}[P^{r,\delta}] - \mathbb{E}[P]\right). \tag{48}$$

Now, using the fact that $\mathbb{E}[\hat{P}_L^{r,\delta}] = \mathbb{E}[P_L^{r,\delta}]$ it is easy to see that

$$\mathbb{E}\left[\left(\hat{P}_L^{r,\delta} - \mathbb{E}[P]\right)^2\right] = \mathrm{Var}(\hat{P}_L^{r,\delta}) + \left(\mathbb{E}[P_L^{r,\delta}] - \mathbb{E}[P^{r,\delta}]\right)^2 + \left(\mathbb{E}[P^{r,\delta}] - \mathbb{E}[P]\right)^2 + 2\left(\mathbb{E}[P_L^{r,\delta}] - \mathbb{E}[P^{r,\delta}]\right)\left(\mathbb{E}[P^{r,\delta}] - \mathbb{E}[P]\right). \tag{49}$$

The desired result then follows from the bound in (47):

$$\mathbb{E}\left[\left(\hat{P}_L^{r,\delta} - \mathbb{E}[P]\right)^2\right] \leq \mathrm{Var}(\hat{P}_L^{r,\delta}) + \left(\mathbb{E}[P_L^{r,\delta}] - \mathbb{E}[P^{r,\delta}]\right)^2 + 4c^2\delta^{2(r+1)} + 4c\delta^{r+1}\left|\mathbb{E}[P_L^{r,\delta}] - \mathbb{E}[P^{r,\delta}]\right| \tag{50}$$

References

[1] R. Draxler, D. Arnold, M. Chino, S. Galmarini, M. Hort, A. Jones, S. Leadbetter, A. Malo, C. Maurer, G. Rolph, K. Saito, R. Servranckx, T. Shimbori, E. Solazzo, G. Wotawa, World Meteorological Organization's model simulations of the radionuclide dispersion and deposition from the Fukushima Daiichi nuclear power plant accident, Journal of Environmental Radioactivity 139 (2015) 172 – 184. doi:10.1016/j.jenvrad.2013.09.014.

[2] H. F. Dacre, A. L. Grant, R. J. Hogan, S. E. Belcher, D. Thomson, B. Devenish, F. Marenco, M. Hort, J. M. Haywood, A. Ansmann, et al., Evaluating the structure and magnitude of the ash plume during the initial phase of the 2010 Eyjafjallajökull eruption using lidar observations and NAME simulations, Journal of Geophysical Research: Atmospheres 116 (D20). doi:10.1029/2011JD015608.

[3] H. N. Webster, D. J. Thomson, B. T. Johnson, I. P. C. Heard, K. Turnbull, F. Marenco, N. I. Kristiansen, J. Dorsey, A. Minikin, B. Weinzierl, U. Schumann, R. S. J. Sparks, S. C. Loughlin, M. C. Hort, S. J. Leadbetter, B. J. Devenish, A. J. Manning, C. S. Witham, J. M. Haywood, B. W. Golding, Operational prediction of ash concentrations in the distal volcanic cloud from the 2010 Eyjafjallajökull eruption, J. Geophys. Res. 117 (D20). doi:10.1029/2011JD016790.

[4] A. Jones, Atmospheric dispersion modelling at the Met Office, Weather 59 (2004) 311–316. doi:10.1256/wea.106.04.

[5] A. Jones, D. Thomson, M. Hort, B. Devenish, The U.K. Met Office's Next-Generation Atmospheric Dispersion Model, NAME III, Air Pollution Modeling and Its Application XVII (2007) 580–589doi:10.1007/978-0-387-68854-1_62.

[6] F. B. Smith, M. J. Clark, The transport and deposition of airborne debris from the Chernobyl nuclear power plant accident with special emphasis on the consequences to the United Kingdom, Meteorological Office Scientific Paper No. 42.

[7] A. L. Redington, R. G. Derwent, C. S. Witham, A. J. Manning, Sensitivity of modelled sulphate and nitrate aerosol to cloud, pH and ammonia emissions, Atmospheric Environment 43 (20) (2009) 3227–3234. doi:10.1016/j.atmosenv.2009.03.041.

[8] S. Heinrich, Multilevel Monte Carlo Methods, in: S. Margenov, J. Waśniewski, P. Yalamov (Eds.), Large-Scale Scientific Computing, Vol. 2179 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2001, pp. 58–67. doi:10.1007/3-540-45346-6_5.

[9] M. B. Giles, Multilevel Monte Carlo Path Simulation, Operations Research 56 (3) (2008) 607–617. doi:10.1287/opre.1070.0496.

[10] M. B. Giles, Multilevel Monte Carlo methods, Acta Numerica 24 (2015) 259. doi:10.1017/S096249291500001X.

[11] M. Giles, L. Szpruch, Multilevel Monte Carlo methods for applications in finance, arXiv preprint arXiv:1212.1377.
URL https://arxiv.org/abs/1212.1377

[12] I. G. Graham, R. Scheichl, E. Ullmann, Mixed finite element analysis of lognormal diffusion and multilevel Monte Carlo methods, Stochastics and Partial Differential Equations Analysis and Computations 4 (1) (2016) 41–75. doi:10.1007/s40072-015-0051-0.

[13] E. H. Müller, R. Scheichl, T. Shardlow, Improving multilevel Monte Carlo for stochastic differential equations with application to the Langevin equation, Proceedings of the Royal Society A 471 (2176). doi:10.1098/rspa.2014.0679.

[14] A. V. Skorokhod, Stochastic Equations for Diffusion Processes in a Bounded Region, Theory of Probability & Its Applications 6 (3) (1961) 264–274. arXiv:http://dx.doi.org/10.1137/1106035, doi:10.1137/1106035.

[15] M. Bernardo, C. Budd, A. R. Champneys, P. Kowalczyk, Piecewise-smooth dynamical systems: theory and applications, Vol. 163, Springer Science & Business Media, 2008.

[16] J. D. Wilson, T. K. Flesch, Flow boundaries in random-flight dispersion models: enforcing the well-mixed condition, Journal of Applied Meteorology 32 (11) (1993) 1695–1707.

[17] H. C. Rodean, Stochastic Lagrangian Models of Turbulent Diffusion, Vol. 26, American Meteorological Society, 1996.

[18] M. B. Giles, T. Nagapetyan, K. Ritter, Multilevel Monte Carlo Approximation of Distribution Functions and Densities, SIAM/ASA Journal on Uncertainty Quantification 3 (1) (2015) 267–295. doi:10.1137/140960086.

[19] P. Kloeden, E. Platen, Numerical Solution of Stochastic Differential Equations, Stochastic Modelling and Applied Probability, Springer Berlin Heidelberg, 2011.
URL https://books.google.co.uk/books?id=BCvtssom1CMC

[20] B. Leimkuhler, C. Matthews, Molecular Dynamics: With Deterministic and Stochastic Numerical Methods, Interdisciplinary Applied Mathematics, Springer International Publishing, 2015.
URL https://books.google.co.uk/books?id=-5oxrgEACAAJ

[21] G. Maruyama, Continuous Markov processes and stochastic equations, Rendiconti del Circolo Matematico di Palermo 4 (1) (1955) 48–90.

[22] N. Bou-Rabee, H. Owhadi, Long-run accuracy of variational integrators in the stochastic context, SIAM Journal on Numerical Analysis 48 (1) (2010) 278–297. doi:10.1137/090758842.

[23] H. Hoel, E. von Schwerin, A. Szepessy, R. Tempone, Adaptive Multilevel Monte Carlo Simulation, in: Numerical Analysis of Multiscale Computations: Proceedings of a Winter Workshop at the Banff International Research Station 2009, Vol. 82, Springer Science & Business Media, 2011, p. 217.

[24] H. Hoel, E. Von Schwerin, A. Szepessy, R. Tempone, Implementation and analysis of an adaptive multilevel Monte Carlo algorithm, Monte Carlo Methods and Applications 20 (1) (2014) 1–41. doi:10.1515/mcma-2013-0014.

[25] M. B. Giles, C. Lester, J. Whittle, Non-nested Adaptive Timesteps in Multilevel Monte Carlo Computations, Springer International Publishing, Cham, 2016, pp. 303–314. doi:10.1007/978-3-319-33507-0_14.

[26] E. H. Mueller, G. Katsiolides, T. Shardlow, MLMCLangevin code (Oct 2016). doi:10.5281/zenodo.161664.

[27] T. Shardlow, Modified equations for stochastic differential equations, BIT Numerical Mathematics 46 (1) (2006) 111–125. doi:10.1007/s10543-005-0041-0.

[28] K. C. Zygalakis, On the existence and the applications of modified equations for stochastic differential equations, SIAM Journal on Scientific Computing 33 (1) (2011) 102–130. doi:10.1137/090762336.

[29] D. J. Thomson, Criteria for the selection of stochastic models of particle trajectories in turbulent flows, Journal of Fluid Mechanics 180 (1987) 529–556. doi:10.1017/S0022112087001940.

[30] A. N. Kolmogorov, The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers, in: Dokl. Akad. Nauk SSSR, Vol. 30, JSTOR, 1941, pp. 301–305.
URL http://www.jstor.org/stable/51980

[31] J. Wilson, E. Yee, N. Ek, R. d'Amours, Lagrangian simulation of wind transport in the urban environment, Quarterly Journal of the Royal Meteorological Society 135 (643) (2009) 1586–1602. doi:10.1002/qj.452.

[32] H. Webster, D. Thomson, N. Morrison, New turbulence profiles for NAME, Met Office Turbulence and Diffusion note 288.

[33] R. Khasminskii, Stochastic stability of differential equations, Vol. 66, Springer Science & Business Media, 2011.

[34] H. Hoel, E. Schwerin, A. Szepessy, R. Tempone, Numerical Analysis of Multiscale Computations: Proceedings of a Winter Workshop at the Banff International Research Station 2009, Springer Berlin Heidelberg, 2012, Ch. Adaptive Multilevel Monte Carlo Simulation, pp. 217–234. doi:10.1007/978-3-642-21943-6_10.

[35] H. Hoel, E. Von Schwerin, A. Szepessy, R. Tempone, Implementation and analysis of an adaptive multilevel Monte Carlo algorithm, Monte Carlo Methods and Applications 20 (1) (2014) 1–41. doi:10.1515/mcma-2013-0014.

[36] W. Fang, M. B. Giles, Adaptive Euler-Maruyama Method for SDEs with Non-globally Lipschitz Drift: Part I, Finite Time Interval, arXiv preprint arXiv:1609.08101.

[37] F. Y. Kuo, R. Scheichl, C. Schwab, I. H. Sloan, E. Ullmann, Multilevel quasi-Monte Carlo methods for lognormal diffusion problems, arXiv preprint arXiv:1507.01090.
URL https://arxiv.org/abs/1507.01090